

# Analyzing and Simulating User Utterance Reformulation in Conversational Recommender Systems

Shuo Zhang\*  
Bloomberg  
London, United Kingdom  
imsure318@gmail.com

Mu-Chun Wang\*  
University of Science and Technology  
of China  
Hefei, China  
amoswang2000@mail.ustc.edu.cn

Krisztian Balog  
University of Stavanger  
Stavanger, Norway  
krisztian.balog@uis.no

## ABSTRACT

User simulation has been a cost-effective technique for evaluating conversational recommender systems. However, building a human-like simulator is still an open challenge. In this work, we focus on how users reformulate their utterances when a conversational agent fails to understand them. First, we perform a user study, involving five conversational agents across different domains, to identify common reformulation types and their transition relationships. A common pattern that emerges is that persistent users would first try to rephrase, then simplify, before giving up. Next, to incorporate the observed reformulation behavior in a user simulator, we introduce the task of reformulation sequence generation: to generate a sequence of reformulated utterances with a given intent (rephrase or simplify). We develop methods by extending transformer models guided by the reformulation type and perform further filtering based on estimated reading difficulty. We demonstrate the effectiveness of our approach using both automatic and human evaluation.

## CCS CONCEPTS

• Information systems → Users and interactive retrieval.

## KEYWORDS

Conversational recommender system, utterance reformulation, user study, user simulation

### ACM Reference Format:

Shuo Zhang, Mu-Chun Wang\*, and Krisztian Balog. 2022. Analyzing and Simulating User Utterance Reformulation in Conversational Recommender Systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*, July 11–15, 2022, Madrid, Spain. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3477495.3531936>

\*The first two authors contributed equally.

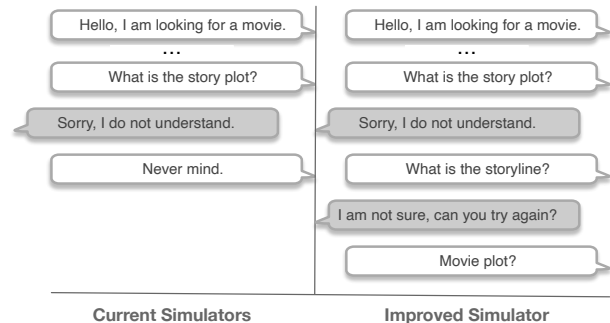
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

SIGIR '22, July 11–15, 2022, Madrid, Spain

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-8732-3/22/07...\$15.00

<https://doi.org/10.1145/3477495.3531936>



**Figure 1: Existing simulators (Left) will give up immediately when the agent fails to understand them. Our improved simulator (Right) makes multiple attempts by reformulating the request.**

## 1 INTRODUCTION

A conversational recommender system (CRS) is a task-oriented dialogue system that supports its users in fulfilling recommendation-related intentions via a multi-turn conversational interaction [19]. A major differentiating factor from traditional recommender systems is that interactions are in natural language [15], which allows for a more natural elicitation of preferences [31] and nuanced feedback on the suggested items [6]. Despite the tremendous advances made in NLP technology in recent years, natural language understanding continues to present a major challenge for CRSs [15, 19].

Progress in this area critically depends on the availability of appropriate evaluation methodology and resources. Therefore, evaluation of conversational recommender systems, in and of itself, is a vitally important research topic [15, 19]. In general, recommender systems can be evaluated along various dimensions, including the quality of the recommendations, the perceived transparency of the process, or the ease-of-use of the system [19]. While the same criteria can be also applied for conversational recommender systems, given their interactive nature, human-computer interaction aspects also need to be considered, related to the efficiency or quality of the conversation itself [19]. End-to-end system evaluation is no longer possible using offline test collections, as due to its dynamic nature the conversation can branch “infinitely” at each user turn. Therefore, offline test collections are limited to a single turn and a predefined set of possible user requests [11]. Annotating entire dialogues using human judges is an alternative, but it is expensive, time-consuming, and does not scale [53]. A promising viable alternative is user simulation, which would provide a scalable and cost-efficient solution to automatic evaluation [4, 5, 15].

User simulation has been employed in the past in the context of task-oriented dialogue systems [2, 16, 37]. There, simulation is primarily used for training reinforcement learning-based dialogue systems [12, 18, 20], and less so for evaluation. Apart from a few notable exceptions [41], those user simulators are relatively simplistic, using pattern-based replies. More recently, the idea of simulation has been revisited for conversational search and recommender systems, targeting more task-specific requirements. Notably, Salle et al. [36] focus on generating answers to clarification questions, while Zhang and Balog [53] develop models to guide recommendation-specific dialogue flows and to ensure the consistency of personal preferences.

A main shortcoming of existing user simulators is that they offer limited robustness when the generated user utterance is not understood by the system. In the absence of a better option, they revert back to either restarting or terminating the dialogue. Our objective is to develop a more human-like user simulator that would not “give up” immediately, if a given utterance was not understood by the system; see Figure 1. The first main research question we ask is: *How do users behave when their requests are not understood by a CRS?* To answer this question, we design a user study and collect data using crowdsourcing on five conversational agents covering different domains, including movies, music, and travel, and annotate the dialogues using an established categorization of intents [53]. We observe a typical behavioural pattern, according to which a persistent user would first try to rephrase, then to simplify their request, before giving up.

Motivated by these findings, we formulate a second main research question: *How to model human-like reformulation behavior?* To answer it, we introduce the task of *reformulation sequence generation*: generating a sequence of different expressions for the same user intent, each with a specific reformulation type (rephrase or simplify). This task resembles the problems of question rewriting in conversational QA [1, 25, 39, 45] and query suggestions in conversational information-seeking tasks [35]. However, reformulation sequence generation differs from question rewriting in major ways: as opposed to generating a single rewrite of the query, it outputs a diverse set of utterances corresponding to different reformulation types in the particular order in which they would be inputted by the user, should the system fail to understand them.

We propose to solve this task using a type-guided transformer, where the target reformulation type is provided as input. While these models can generate different reformulations of the same utterance, not all are appropriate in a sense that they do not simplify the language. Therefore, we additionally investigate a way to estimate reading difficulty and reject utterances that are deemed linguistically unacceptable. This reformulation behavior can then be incorporated in the user simulation architecture via the response generation and natural language generation components.

We evaluate our approach on the component-level, using both automatic and human evaluation. Additionally, we perform contextual evaluation, by substituting human reformulation sequences with simulated ones in actual dialogues, and measuring to what extent human judges are able to distinguish between the two. We find that utterance reformulation contributes to making simulation more human-like, and a combination of type and reading difficulty

can improve the utterance reformulation performance consistently across domains, methods, and evaluation measures.

In summary, the novel contributions of this paper are as follows.

- We conduct a user study to understand the traits of human utterance reformulation across five agents and three domains.
- We introduce the novel task of simulating utterance reformulation and explore how different transformer-based models can be used for addressing this task.
- We perform rigorous evaluation on the component-level as well as in an actual dialogue context.
- We construct a data collection consisting of dialogues by 150 human users, expert annotations for each dialogue turn, and 1.8k reformulated conversational utterances that may be used for training more realistic user simulators.

Resources are shared in the repository accompanying this paper: <https://github.com/iai-group/sigir2022-usersim>.

## 2 RELATED WORK

This study lies in the intersection of conversational information access, evaluation, and utterance rewriting.

### 2.1 Conversational Information Access

*Conversational information access* refers to a subset of task-oriented dialogue systems aimed at supporting users in fulfilling information-seeking goals via multi-turn conversations [53]. It includes the tasks of conversational search, QA, and conversational recommendations [52]. Commercial products, like Amazon’s Alexa, Google Assistant, Apple’s Siri, Microsoft’s Xiaolce, are assisting users in undertaking specific tasks in restricted domains, like restaurant reservations and music recommendation. Recent studies have made progress in improving specific subtasks, including response generation, suggesting relevant queries [35], asking clarifying question [1], predicting user intent, and user preference modeling [53]. Azzopardi et al. [3] classify these supported actions and interactions identified in [32] as *query formulation*, *set retrieval*, and *mixed-initiative*. There also exists non-task-oriented dialogues systems, also known as chatbots, which support unstructured human-human interactions like chit-chat [15]. However, we concentrate on the conversational recommender system in this work [19].

### 2.2 Evaluation

Compared to traditional recommender systems, the evaluation of conversational recommender systems is focused more on the human-computer interaction aspects [14]. Jannach et al. [19] identify four main evaluation dimensions for conversational recommenders: (1) effectiveness of task support, (2) efficiency of task support, (3) quality of the conversation and usage, and (4) efficiency of the subtask. Each dimension considers different measurements on the component level and system level. For example, component measures for NLU include the accuracy, precision, and recall, while for NLG word-overlap based metrics are commonly employed [7, 26]. System-level metrics include reward and hit rate [26, 28]. Due to the dynamic nature of interactions, traditional (offline) evaluation methodology [9] is not suitable; alternative means of evaluation include user simulation and user studies.

*User simulation* has been widely leveraged in the past for training the dialogue state tracking component of conversational agents using reinforcement learning algorithms, either via agenda-based or model-based simulation [19]. The highly interactive nature of conversational information access systems has also sparked renewed interest in evaluation using user simulation within the IR community [4, 5, 23, 36, 38, 53]. Recently, Zhang and Balog [53] proposed a general framework for evaluating conversational recommender systems using user simulation. Their agenda-based simulator comprises NLU, NLG, and response generation components. However, building a human-like simulator is still an open challenge [4]. To bridge the gap, Salle et al. [36] focus on behavioral factors like cooperativeness and patience to build a more human-like simulator for information-seeking conversations. Sekulic et al. [38] take this further by enabling simulated users to ask clarifying questions in a mixed-initiative setting. Sun et al. [40] studied how to simulate user satisfaction in a human-like way for task-oriented tasks. In this work, we concentrate on improving human-likeness in terms of how users would reformulate their utterances in case they are not understood by the system.

The use of *user studies* is a well-established human evaluation approach to understand user behavior in interactive tasks, e.g., search result selection [46], query formulation [43], and user engagement with virtual assistants [29, 44, 47, 48]. User studies may be categorized as lab, in-field, and crowdsourcing experiments [12]. For example, Trippas et al. [43] conduct a lab-based user study to observe the characteristic of spoken interactions, such as how to initialize the query and reformulate them, and find that spoken conversational search is more complex and interactive than traditional search. In [44] an in-field survey is performed with 401 respondents on daily tasks and activities with intelligent assistants in a work setting, which helps to understand the role of these assistants in supporting people engaged in work tasks. The MSDialog dataset [29] contains conversation data collected using crowdsourcing. Qu et al. [29] find that short conversation cues such as “Okay” and “Yup” take place frequently in guided tasks with virtual assistants. Vtyurina and Fourney [47] carry out a Wizard-of-Oz study to study the role of implicit cues and conclude that a conversational system that is able to recognize and act upon such cues would enable its users to converse using a more natural style of language.

## 2.3 Utterance Reformulation

*Utterance reformulation* or *utterance rewriting* is related to the well-studied problem of query reformulation in IR [34], and has been explored in many conversational QA tasks, e.g., for ambiguous questions [45], incomplete utterances [25], multi-turn dialogue modeling [39], and query clarification [1]. Specifically, Vakulenko et al. [45] propose to reformulate ambiguous questions, whose reference information is only introduced earlier in the conversation, into implicit ones based on the context for the task of conversational QA. Liu et al. [25] consider incomplete utterance rewriting as a semantic segmentation task, casting it as a word-level edit matrix construction problem. To improve the performance of multi-turn dialogue modeling, Su et al. [39] simplify the multi-turn into a single-turn task as a pre-process through utterance rewriting, which includes recovering all co-referred and omitted information.

Based on a purpose-built 200k multi-turn conversational dataset, a pointer-generator model is trained as a rewriter for this task, which can be integrated into online chatbots. Aliannejadi et al. [1] propose to ask clarifying questions when users fail to formulate their complex information needs in a single query. Similarly, Rosset et al. [35] suggest useful questions in conversational search to improve the users’ engagement. To enhance the diversity and quality of system responses, Lippe et al. [24] combine two commonly used response generation methods, template-based and corpus-based, by paraphrasing template-based responses. In this paper, we focus on a particular reason why reformulation happens: when the conversational agent fails to understand the user utterance. We aim to model this reformulation behavior as a sub-component of the user simulator architecture.

## 3 ANALYZING UTTERANCE REFORMULATIONS

Most conversational recommender systems are still in an early stage of development, and many questions have not been well studied, including user behavior on utterance formulation [43]. We carry out a user study to answer the following question: *How do users behave when their requests are not understood by a conversational recommender system?* To ensure that our observations are based on realistic user behavior, we need users to interact with an actual agent, as opposed to Wizard-of-Oz style studies, with an idealized agent behavior. We design a crowdsourcing experiment on Amazon MTurk<sup>1</sup> that involves 150 participants engaging with five conversational agents. Based on the collected dialogues, we find that utterance rephrasing and simplification are the two most frequent types of behavior, and they take place equally frequently when interacting with a conversational recommender.

### 3.1 Experimental Design and Participants

**3.1.1 Objectives and Design Decisions.** Our main goal with this experiment is to understand how humans reformulate utterances when a conversational agent fails to understand them. To ensure that our findings are not specific to a particular agent, but are generalizable observations, we consider a number of different agents.

Another design decision we make is instructing the study subjects to keep communicating with the conversational system until they get a satisfactory recommendation. This is to mimic the behavior of a persistent user, which clearly introduces some bias. However, we consider this behavior as corresponding to maximum persistence, whereas the other extreme would be giving up immediately. The degree of user persistence can be adjusted and incorporated as a parameter in the simulator (see, e.g., [36]).

**3.1.2 Conversational Agents.** We carefully choose five existing third-party conversational agents that each of the systems should at least support the query formulation and set retrieval functionalities (cf. Sect.2.1). Three of these are for movie recommendation, one for music, and one for the travel domain.

- **A1:** *And chill*<sup>2</sup> is a movie theme chatbot that interacts on Facebook and provides Netflix recommendations. After answering a

<sup>1</sup><https://www.mturk.com/>

<sup>2</sup><http://www.andchill.io/>

few questions such as a liked movie and the reason why liking it, the agent sends movie recommendations based on users' preferences.

- **A2: Jarvis [53]** is a movie recommender system that is based on the Plato Research Dialogue System.<sup>3</sup> This agent can answer questions about movies such as directions, and ratings, and provide recommendations based on genres. It also solicits feedback on movies the user has watched.
- **A3: IAI Movie Bot [17]** is an open-source movie recommender system with a multi-modal chat interface. It focuses on modeling users' preferences dynamically via multi-turn dialogue.
- **A4: Eddy Travel Bot** is a travel assistant that can help users organizing their trip on Telegram.<sup>4</sup> Based on the travel destination, the agent gives recommendations, including flights, hotels, and restaurants.
- **A5: VKM Bot** is a music bot on Telegram that helps users find, listen, and download songs.<sup>5</sup> The agent recommends a list of songs based on a song or artist name provided by the user.

A2 and A3 are two similar agents in terms of supported functionalities. However, they are built on different dialogue frameworks. Besides, A2 is a conversational agent without "keyboard guidance" (i.e., buttons with suggestions) for users while A3 takes more initiative and guides users by offering them a set of options on buttons. We observe that users reformulate less with keyboard-based agents.

**3.1.3 Crowdsourcing Experiment.** The user study comprises of two parts: demographic information collection and dialogue collection. To learn how user behavior changes across user groups, we obtain demographic information by designing a small questionnaire, where we ask about the participants' background, such as age, gender, and education, and whether they have experience with conversational agents. We invite crowd workers on Amazon MTurk as participants and provide incentive payment for each task (\$2.5 for each task). Each participant is instructed to interact with a conversational recommender system either on Facebook or Telegram. They need to engage with the agent to find a liked movie, song, or travel plan. They were given the following instructions:<sup>6</sup>

*In this task, you are going to interact with a chatbot that will help you find movies that you might like. Your task is to engage in a conversation with this bot until you receive a recommendation for an item that you haven't seen before but you'd be interested in watching. Note that the bot may not understand everything you say, in which case you'll need to try again by stating your request somewhat differently.*

For each of the agents, we invite around 60 participants to complete the task. We manually choose 30 questionnaires per agent after quality control<sup>7</sup> and select a total of  $5 \times 30 = 150$  dialogues (that is, by 150 unique users), which will be analyzed. Of the 150 participants, 42% are under 30 years old, 43.3% are between 30 and 40, and 14.7%

are older than 40. In terms of gender, 46% are female and 54% are male. 38.7% of the annotators do not have higher education, 44.7% have bachelor or college study, 14.7% have a master degree, and 2% have doctoral education. Most of the annotators have prior experiences with using conversational chatbots such as Google Assistance, Siri, and Amazon Alexa, for finding restaurants, music, shopping, sports, or movies. However, 37% have never asked for recommendations.

## 3.2 Understanding Logged Dialogues

To distill the utterance reformulation behavior from the logged dialogues, expert annotations were performed (by two of the paper's authors) on the exchange level along three dimensions: intent, slot, and reformulation.<sup>8</sup> In addition to facilitating user behavior analysis, these annotations also enable the use of the collected dialogues as training data for user simulation in the future.

**3.2.1 Intent Annotation.** Intent reflects the action or interaction of the current utterance. We take the actions and interactions identified in [53] as intents, which are classified more broadly in [32] as *query formulation*, *set retrieval*, and *mixed-initiative*. In the phase of query formulation, the system will *elicit* the user to *reveal* information need, which is further classified as *disclose*, *non-disclose*, *revise*, *refine*, and *expand*. *Set retrieval* is to display search and recommendation results, and it includes *inquire* (*list*, *compare*, *subset*, *similar*) and *navigate* (*repeat*, *back*, *more*, *note*, *complete*). *Mixed initiative* include *interrupt* and *interrogate*.

**3.2.2 Slot Annotation.** In conversational recommendations, a slot refers to the variable for categorizing and interpreting the utterances in both the preferences and recommendations [15]. For example, a movie mention or genre in *disclose* reflects the user's tastes and preferences. We annotate specific items (movie, location, restaurant, hotel, song, musician) and movie genres.

**3.2.3 Reformulation Annotation.** We assume that reformulation is likely to happen when the agent fails to understand or misunderstands the user. Based on the logged dialogues, we develop a new coding scheme, with a particular focus on capturing reformulation behavior; see the possible types and examples in Table 1. We annotate our dialogue collection with these reformulation types.

According to the reformulation types listed in Table 1, *Start/restart* indicates the start of articulating an information need (i.e., set of requirements). In the example, the agent fails to recommend more items, thus the user restarts the dialogue by disclosing extra preferences. When the agent fails to understand the user or simply gives an irrelevant reply, the user might *repeat* the last utterance without significant change. *Repeat* can take place by repeating the last utterance using different words/expression (about the same complexity), namely *Repeat/rephrase*, or repeating the last utterance a simpler/shorter words/expression with reduced complexity (*Repeat/simplify*). When the agent does not completely understand the user, the user may clarify or refine the constraints expressed in their last utterance, denoted as *Clarify/refine*. Under similar circumstances, the user might change the constraints (adding or removing) expressed in the last user utterance, which is defined as *Change*.

<sup>3</sup><https://github.com/uber-archive/plato-research-dialogue-system>

<sup>4</sup>[https://t.me/eddytravels\\_bot](https://t.me/eddytravels_bot)

<sup>5</sup>[https://telegram.me/vkm\\_bot](https://telegram.me/vkm_bot)

<sup>6</sup>The instructions were piloted on a smaller-scale lab-based study first (involving five undergraduate computer science students) to ensure that they are understandable.

<sup>7</sup>We exclude dialogues with very few turns, failure in asking for recommendations and repetitive records.

<sup>8</sup>Disagreements during the annotation process were resolved via discussion.

**Table 1: Types of reformulation behavior in conversational recommender systems, where the agent fail to understand/satisfy the user. Each example lists three consecutive utterances, where two are from the user (U) and one by the agent (A).**

Type	Definition	Example
<b>Start/restart</b>	Start of articulating an information need (i.e., set of requirements)	U: ..., but I'm looking for a movie I've never seen. A: That's all I got at the moment. U: Dumb and Dumber is my favorite comedy...
<b>Repeat</b>	Repeat the last utterance without significant change	U: I am looking for amazon premium movie. A: I'm pretty solid on a bunch of things so far. U: I am looking for amazon premium movie.
<b>Repeat/rephrase</b>	Repeat the last utterance using different words/expression (about the same complexity)	U: No, I'm looking for a restaurant. A: Sorry, I didn't get that. Can you rephrase? U: What restaurants do you know of in Fort Lauderdale?
<b>Repeat/simplify</b>	Repeat the last utterance a simpler/shorter words/expression (reduced complexity)	U: Never mind. Can you book me a taxi from the airport? A: Sorry, I didn't get that. Can you rephrase? U: A cab.
<b>Clarify/refine</b>	Clarify or refine the constraints expressed in the last user utterance	U: Something light-hearted. A: Can you tell me about a different one you like? U: I want a funny comedy that is light hearted.
<b>Change</b>	Change the constraints (adding or removing) expressed in the last user utterance	U: I don't think either of those are romantic comedies. A: I think you should give Balls Of Fury a shot! U: Okay, let's try a movie that is similar to Speed.
<b>Stop</b>	End the session	U: I give it a year looks really good, why did you recommend that one? A: Great! I'm really excited. These movies are going to be phenomenal! U: Thanks again, I think I watch I Give it a Year.

**Table 2: Agent intents that led to reformulations with relative frequencies (Ratio) and standard dev. across agents ( $\sigma$ ).**

Agent Intent	Ratio	$\sigma$	Agent Intent	Ratio	$\sigma$
Failed	0.62	0.0021	Similar	0.02	0.0013
Suggest	0.19	0.0015	Repeat	0.02	0.0021
Elicit	0.06	0.0017	Non-disclose	0.02	0.0018
Extract	0.03	0.0017	End-disclose	0.01	0.0015
List	0.03	0.0015	Clarify	0	0

Finally, the user might *End* the session once the information needs have been met, though the agent response is not as expected.

### 3.3 Analyzing Reformulations

Next, we analyze the annotated dialogues with the aim to gain an understanding of humans' reformulation behavior.

**(RQ1): When do reformulations occur?** We focus on the last agent utterances that led to users' reformulations, making use of the intent annotations in Sect. 3.2.1. Table 2 displays the ratio of these intents. We find that reformulation takes place mostly when encountering *Failed* replies, i.e., when the agent fails to understand the user. But it also appears with other system intents, such as *Suggest* and *Elicit*. There are some indications of differences when considering demographic information, such as gender and age. For example, when age is over 40, the ratio of *Failed* intent has dropped to 45%, while *Elicit* has raised to 23%. The reason behind this might be that users' desire to respond changes with age. Younger users are more proactive in interacting with the agents when the agent failed.

Older users, on the other hand, seem to be more patient with other intents, e.g., witnessed by a high ratio for reformulation in *Elicit*. We additionally investigate the ratios for each conversational agent and find that their distributions are similar to each other, witnessed by the low standard deviations. Besides, the distributions of agent intent are similar across different ages, gender, and experiences. We note that these observations are indicative but not conclusive, as our data sample is limited in size. A larger-scale analysis of demographic differences is left for future work.

**(RQ2): Do people who have experience with conversational agents show different reformulation behavior from those that do not have that?** For each intent, we calculate the distributions of reformulation types for users both with and without experience with conversational recommendations. For each pair of distributions, we first take the Levene test [8] to examine the homogeneity of variance between them for each reformulation type. After this, we perform a t-test based on the mean values and find that the all the p-values are much larger than 0.1, meaning these two groups do not have differently significant behavior with regards to these intents. These statistics are based on aggregated data across all agents, but we find that the same observations also hold for individual agents.

**(RQ3): Can reformulation patterns be identified?** The sequential reformulations reveal how reformulation behavior changes when encountering continuous dissatisfying system replies. To answer RQ3, we calculate the transition probabilities of the reformulation states. We split the logged dialogues into dialogue pieces by intents, where each reformulation is a state. We tracked the utterance by




Figure 2: Reformulation sequence generation architecture.

Table 3: Top reformulation patterns in our dataset.

Intent	Ratio	Count
Rephrase-Simplify	0.13	80
Rephrase-Refine	0.10	63
Simplify-Rephrase	0.08	53
Refine-Simplify	0.03	20
Simplify-Refine	0.03	19
Refine-Rephrase	0.03	18

intent and slot; if the slot and intent have changed, we consider it as another dialogue piece. One dialogue piece might occur in multiple states as the state changes upon the change of reformulation type. For example, the user might reformulate differently for the same intent (in the same dialogue piece). We construct the transition matrix  $M$  based on the maximum likelihood estimation method [42]. In particular, we consider the state transition sequence of each agent as a Markov process, and further maximize the likelihood for each element:

$$m_{ij} = p(t_r = j | t_{r-1} = i) = \frac{\sum_r \mathbb{1}(t_{r-1} = i \wedge t_r = j)}{\sum_r \mathbb{1}(t_{r-1} = i)},$$

where  $t$  refers to the reformulation type, and  $r$  means the current conversation turn. The frequencies of top transition patterns are reported as in Table 3. For consecutive reformulations, we observe a typical pattern where a persistent user would first try to *rephrase*, then *simplify* before giving up.

Additionally, we ask: How do the type of reformulations change during the course of dialogues? In Fig. 3, we show the distribution of various reformulation types across different stages (binned by turns of 5) on the bar plot (top), while the line plot (bottom) shows the absolute volume (i.e., total number of turns where reformulations occurred). We observe, again, that rephrasing tends to happen earlier in dialogues than simplification.

## 4 SIMULATING UTTERANCE REFORMULATIONS

In Section 3, we have analyzed user behavior in terms of utterance reformulations. Next, we address the question: *How to incorporate human-like reformulation behavior into a state-of-the-art user simulator?* Specifically, we propose the reformulation sequence




Figure 3: Reformulation types over dialogue turns.

generation task (cf. Sect 4.1) to output a reformulation sequence based on a given utterance and reformulation type (rephrase or simplify), using type-guided text generation transformers (Sect. 4.2). To ensure that the reformulations indeed achieve their intended goal, we further investigate ways to estimate the reading difficulty of the generated utterances (Sect. 4.3).

### 4.1 Reformulation Sequence Generation

Users might attempt to reformulate the utterance when their requests are not understood. We describe these sequences of utterances resulted from reformulations as a reformulation sequence,  $S^R$ . Each  $S^R$  includes at least two user utterances (conversation turns) with the same user intent. To simulate a reformulation sequence, we propose the reformulation sequence generation (RSG) task. Given a user utterance  $u_0$ , it aims to generate a sequence of reformulations:

$$RSG(u_0) = S^R = \{u_1, \dots, u_n\}, \quad (1)$$

where  $u$  denotes a user utterance. For example, if  $u_0$  is “I am into a movie like a tale of two cities”, the user might *rephrase* it to  $u_1$  “movies similar to a tale of two cities” when the system fails to understand the preference disclosure. If failure continues, the user might *repeat*  $u_1$  in  $u_2$  as “movies similar to a tale of two cities” or even *simplify*  $u_2$  to  $u_3$  “a tale of two cities”. Figure 2 illustrates the architecture we devise for this task. As a result, the utterances in Eq. 1 change

sequentially and form a sequence of reformulation types:

$$t = \{t_1, t_2, \dots, t_n\}, \quad (2)$$

where  $t_n$  denotes the reformulation type from  $u_{n-1}$  to  $u_n$ . To model these transitions, we leverage the reformulation types identified in Table 1, and leverage the transition matrix  $M$  in Section 3 to determine the reformulation type  $t$ . Specifically, we randomly select the first reformulation type  $t_1$  by the uniform distribution, and we further use it as a condition to generate the first reformulated utterance  $u_1$ . E.g., given  $u_0$  “*Something light-hearted*,” we select *refine* as the first reformulation type  $t_1$ , and further use it as a condition to generate  $u_1$  “*I want a funny comedy that is light hearted*.” After the first generation, we update the distribution of reformulation type (“Dist update” in Fig. 2) by following the transition probability:

$$\tilde{z}_r = \tilde{z}_r | \tilde{z}_{r-1} = M \cdot \tilde{z}_{r-1}, \quad (3)$$

where  $M$  represents the transition matrix, and the vector  $\tilde{z}_r$  represents the probability distribution of reformulation types at step  $r$ . We update the distribution by calculating the dot product of the matrix and previous distribution vector. After getting a new distribution of reformulation types, we select  $t_2$  and use it as a condition to generate  $u_2$ . We then repeat this procedure in Fig. 2 to generate a sequence of reformulated utterances as output. We consider the sequence as a set of candidates which can be used to respond to the agent in simulation, if it fails to give a satisfactory reply.

## 4.2 Type-guided Reformulation Utterance Generation

In Section 4.1, we formulated our task as a conditional sequence to sequence task. In each iteration, given  $u_0$  and next reformulation type (which might be *repeat*, *rephrase*, *simplify*, *repeat*, or *restart*), the model generates the next reformulated utterance via type-guided reformulation generation methods.

We consider different transformer models, both a pre-trained language model and sequence-to-sequence models for addressing this task. To train a language model, we concatenate the reformulation type, utterance, and reformulated utterance and further use masks to identify them. We consider *GPT2* [30] and use top-k sampling for the decoding algorithm, with  $k$  set to 0.5. To train a sequence-to-sequence model, we consider both the utterance and the reformulation type  $t$  as input, further concatenate them (separated by a special  $\langle /s \rangle$  token), and tokenize them as an input tensor. For the output, we tokenize the reformulated utterance as an output tensor. The sequence-to-sequence models we consider for comparison are *BERT-to-BERT* [13] (denoted as *BERT*), *BART* [21], and *Text-to-Text Transfer Transformer (T5)* [33].

## 4.3 Filtering based on Reading Difficulty

One of the objectives of reformulation is to lower the reading difficulty for the generated utterance. To avoid generating less understandable or linguistically more complex utterances than the original one, we include a module for estimating the reading difficulty in our simulation. It classifies each generated utterance as acceptable or not. Only the acceptable ones will be considered as candidates, i.e., the generated utterances with the  $\sqrt$  label in Fig. 2.

There has been a long history of research on language reading difficulty [10, 49] and neural language understanding models advance the development recently [13, 50]. As the first step towards this investigation in user simulation, we leverage a BERT-based classifier [13] fine-tuned on the CoLA dataset [50]. This dataset is annotated based on whether the language is linguistically acceptable or not. Three types of sentences, which are *morphological anomalies*, *syntactic anomalies*, and *semantic anomalies* are labeled as unacceptable. Taking an example for our task, the binary classifier will label the utterance “*I want to watch a good rating*” as unacceptable with unavailable meanings since it is considered as a morphological anomaly. To train the BERT classifier, we follow the exact same experimental setting as in [13].

## 5 EXPERIMENTAL SETUP

### 5.1 Constructing Training Collection

Recall that we have collected data from five conversational agents on three domains. We construct training data for the utterance generation task by taking these logged dialogues from the user study and further enrich them via another crowdsourcing task.

**5.1.1 Logged Dialogues.** To integrate the logged dialogues from the user study, we consider user utterances labeled with a *reformulation* type (cf. Sect. 3.2.3) and further back-trace them to the previous user utterance as the *original utterance*. The original utterance, reformulation type, and reformulated utterance form a triad  $\langle u_0, t, u_r \rangle$ . We consider the triad records from all the agents except *Music*, since there are only 5 reformulation sequences in all 30 logged dialogues for the Music domain. For the other 4 agents, we combine the 30 dialogues per agent as our initial dataset.

**5.1.2 Crowdsourcing.** To enrich the diversity of the initial dataset, we conduct an utterance rewriting crowdsourcing task on Amazon MTurk. We present workers with a scenario that a conversational agent fails to understand the request when recommending movies or travel plans. We instruct the annotators to rewrite an utterance by reformulating, refining, or simplifying it. We sample 208 triads from Sect. 5.1.1 when the system fails to understand the requests. The original reformulation sequence serves as the input to simulate the scenario and enable each rewriting task. We invite nine annotators per utterance to rewrite it based on the intent, and harvest about 1.8k rewritten utterances by 265 workers. These collected utterances are integrated with the logged training data by domain, intent, and reformulation type.

### 5.2 Experimental Methods

**5.2.1 Implementation Details.** We use the Huggingface [51] implementations of BERT-to-BERT, Bart, T5, and GPT2. For the BERT-to-BERT model, we use the Encoder-Decoder Models frame with both initialized with the “bert-uncased” model. To finetune Bart, T5, and GPT2, we use the Conditional Generation model on Huggingface, and “Facebook/bart-large,” “t5,” and “gpt2” pre-trained models.

To examine the generalization capabilities of these models, we leverage the training data in three ways: only *movie* domain, only *travel* domain, and *hybrid*. The *hybrid* mode is meant to represent a domain-independent approach by combining the *travel* and *movie* datasets. We only consider *refine*, *rephrase*, *simplify*, *repeat*, and

**Table 4: Automatic evaluation results. We test significance on the following pairs of methods: with type-guided vs. without, T5 vs. BART, BART vs. BERT, and BERT vs. GPT2. ‡ denotes significance at the 0.005 level.**

Method	Movie				Travel				Hybrid			
	Rouge-1	Rouge-2	Rouge-L	BLEU	Rouge-1	Rouge-2	Rouge-L	BLEU	Rouge-1	Rouge-2	Rouge-L	BLEU
GPT2	0.112	0.031	0.102	0.058	0.056	0.017	0.061	0.049	0.118	0.005	0.092	0.049
GPT2-T	0.113‡	0.031‡	0.106‡	0.072‡	0.053‡	0.015‡	0.065‡	0.072‡	0.125‡	0.028‡	0.137‡	0.055‡
GPT2-R	0.072	0.021	0.096	0.037	0.056	0.015	0.102	0.042	0.106	0.017	0.124	0.105
GPT2-TR	<b>0.086‡</b>	<b>0.032‡</b>	<b>0.132‡</b>	<b>0.128‡</b>	<b>0.152‡</b>	<b>0.031‡</b>	<b>0.127‡</b>	<b>0.073‡</b>	<b>0.182‡</b>	<b>0.032‡</b>	<b>0.173‡</b>	<b>0.156‡</b>
BERT	0.228	0.06	0.221	0.113	0.110	0.023	0.121	0.092	0.260	0.010	0.192	0.112
BERT-T	0.220‡	0.060‡	0.213‡	0.121‡	0.121‡	0.032‡	0.132‡	0.141‡	0.265‡	0.054‡	0.262‡	0.109‡
BERT-R	0.223	0.09	0.193	0.132	0.132	0.038	0.231	0.086	0.243	0.032	0.234	0.213
BERT-TR	<b>0.324‡</b>	<b>0.092‡</b>	<b>0.312‡</b>	<b>0.132‡</b>	<b>0.213‡</b>	<b>0.065‡</b>	<b>0.231‡</b>	<b>0.153‡</b>	<b>0.321‡</b>	<b>0.072‡</b>	<b>0.342‡</b>	<b>0.327‡</b>
BART	0.281	0.108	0.279	2.229	0.281	0.094	0.281	1.572	0.301	0.120	0.302	2.562
BART-T	0.504‡	0.123‡	0.501‡	5.308‡	0.422‡	0.102‡	0.419‡	5.286‡	0.519‡	0.175‡	0.516‡	8.214‡
BART-R	0.343	<b>0.154</b>	0.328	3.326	0.356	<b>0.158</b>	0.312	2.359	0.362	0.163	0.396	3.382
BART-TR	<b>0.521‡</b>	0.124‡	<b>0.517‡</b>	<b>6.532‡</b>	<b>0.437‡</b>	0.100‡	<b>0.433‡</b>	<b>6.473‡</b>	<b>0.585‡</b>	<b>0.285‡</b>	<b>0.532‡</b>	<b>9.481‡</b>
T5	0.410	0.243	0.410	3.472	0.369	0.262	0.367	2.740	0.400	0.245	0.399	4.739
T5-T	0.742‡	0.320‡	0.739‡	19.031‡	0.741‡	<b>0.297‡</b>	0.738‡	21.485‡	0.783‡	0.340‡	0.782‡	20.987‡
T5-R	0.401	0.261	0.399	3.578	0.365	0.240	0.340	3.253	0.420	0.270	0.430	6.242
T5-TR	<b>0.766‡</b>	<b>0.321‡</b>	<b>0.740‡</b>	<b>20.739‡</b>	<b>0.755‡</b>	0.294‡	<b>0.753‡</b>	<b>23.242‡</b>	<b>0.792‡</b>	<b>0.350‡</b>	<b>0.785‡</b>	<b>22.487‡</b>

*restart* as the reformulation types in the experiments, as *stop* and *change* would change the intent of original utterances. We leave the reformulation task with changed intent for future work. Each original utterance  $u_0$  has more than one reformulated candidates under the same reformulation type. We parallelize the records and feed them into the model in the triad data format as mentioned in Sect. 5.1. The training process is through 5-fold cross-validation. We set the batch size as 10 and initialize the model with default parameters. To test the model, we follow the sequence generation procedure in Fig. 2. We set the sequence length as 3, which is the average length of reformulation sequence in the logged dialogues. While updating the reformulation type from the distribution, it will retry if *change* or *stop* is encountered.

**5.2.2 Utterance Generation Variants.** We consider four variants of each utterance generation method. We use **Model** to represent any text generation model we select in Sect. 4.2.

- **[Model]**: The vanilla ones are without reformulation type as input and reading difficulty filter (first row of each block in Table 4).
- **[Model]-T**: The second rows of each block in Table 4 are the ones trained with guided reformulation type.
- **[Model]-R**: The third rows are the methods considering the component of reading difficulty filter.<sup>9</sup>
- **[Model]-TR**: The last rows consider both type and reading difficulty.

## 6 EVALUATION AND ANALYSIS

We evaluate the approaches devised for our novel simulation task (reformulation sequence generation) using automatic evaluation (Sect. 6.1), expert human evaluation (Sect. 6.2), and crowdsourcing human evaluation (Sect. 6.3).

<sup>9</sup>When adding the reading difficulty filter on the BERT-to-BERT model, we relax the boundary of detecting the non-linguistic utterance, since the utterance generated from the BERT-to-BERT model mostly can't pass the linguistic test

### 6.1 Automatic Evaluation with NLP Metrics

For automatic evaluation we use ROUGE [22] and BLEU [27] as our evaluation metrics. Both metrics are often used on text summarization tasks. To guarantee the stability of experimental results, we generate the reformulation sequence five times. We further filter the generated sequences via the utterance reading difficulty component to exclude hard-to-read utterances. We calculate the metrics using micro-averaging and further average the score across each sequence. All the results are reported in Table 4. When evaluating the effect of type guiding ([Method]-T), we find that most of the type-guided methods outperform the vanilla methods (without type-guiding) significantly ([Method]). We also note the performance of methods without type-guiding is not robust (it can vary a lot). Among the three domains, the type-guided methods for hybrid perform comparably with the single domain method, indicates the generalization capabilities of the combined model.

Filtering based on utterance reading difficulty ([Method]-R) can improve performance in some cases. However, it can enhance performance significantly and consistently when combined with type guidance. Indeed, the [Method]-TR achieves the best performance in almost all cases. Therefore, we conclude that type-guiding and reading difficulty estimation often help on their own, but not consistently across all methods and domains. However, combining the two always works. The improvements are substantial and significant.

In terms of transformer selection, we find that the sequence-to-sequence models significantly outperform the language model (GPT2) on all measures. Amongst the pre-trained sequence-to-sequence models, T5 performs the best.

### 6.2 Expert Human Evaluation

Next, we employ human experts (two of the authors) to evaluate the syntax and semantic quality of the generated reformulations. Specifically, we evaluate three aspects: *Grammar*, *Reading Difficulty*, and *Type Accuracy*. For *movie* and *hybrid*, we randomly sample 200 reformulation sequences; for *travel* we use all 150 reformulation



**Table 5: Side-by-side comparison results, with human evaluators guessing which of two dialogs within a given domain was performed by a simulated user (Win) vs. a real one (Loss); a Tie is given when the evaluator could not decide.**

	Movie			Travel			All		
	Win	Lose	Tie	Win	Lose	Tie	Win	Lose	Tie
<b>BART</b>	<b>39 (27%)</b>	53 (36%)	55 (37%)	<b>27 (33%)</b>	29 (36%)	25 (31%)	66 (29%)	82 (36%)	80 (35%)
<b>BART-TR</b>	34 (23%)	35 (24%)	78 (53%)	24 (30%)	25 (31%)	32 (39%)	58 (25%)	60 (26%)	110 (49%)
<b>T5</b>	36 (24%)	53 (36%)	58 (40%)	16 (20%)	35 (43%)	30 (37%)	52 (22%)	88 (39%)	88 (39%)
<b>T5-TR</b>	<b>49 (33%)</b>	45 (31%)	53 (36%)	<b>27 (33%)</b>	19 (23%)	35 (44%)	<b>76 (33%)</b>	64 (28%)	88 (39%)

**Table 6: Expert human evaluation results. The columns are Grammar, Difficulty, and Type. Lower D values correspond to less difficult text (which is desired).**

	Movie			Travel			Hybrid		
	G	D	T	G	D	T	G	D	T
<b>BART</b>	0.90	1.02	0.14	<b>1.00</b>	1.00	0.15	0.94	1.00	0.14
<b>BART-T</b>	0.98	0.97	0.48	0.97	0.86	0.41	0.98	0.83	0.55
<b>BART-R</b>	<b>1.00</b>	0.87	0.23	<b>1.00</b>	0.90	0.18	0.98	0.88	0.22
<b>BART-TR</b>	<b>1.00</b>	<b>0.75</b>	<b>0.65</b>	0.98	<b>0.68</b>	<b>0.64</b>	<b>1.00</b>	<b>0.73</b>	<b>0.72</b>
<b>T5</b>	<b>1.000</b>	0.95	0.22	<b>1.00</b>	0.95	0.14	0.96	0.93	0.16
<b>T5-T</b>	0.97	0.87	0.72	<b>1.00</b>	0.81	0.82	0.98	0.88	0.72
<b>T5-R</b>	<b>1.00</b>	0.82	0.34	<b>1.00</b>	0.85	0.25	0.98	0.81	0.27
<b>T5-TR</b>	<b>0.98</b>	<b>0.69</b>	<b>0.78</b>	<b>1.00</b>	<b>0.73</b>	<b>0.88</b>	<b>0.98</b>	<b>0.72</b>	<b>0.82</b>

sequences. We only evaluate the BART and T5 models as the majority of the utterances generated by the BERT-to-BERT models fail to pass our reading difficulty filter. The annotations are agreed by the authors and the results are reported in Table 6.

Our observations are as follows. The *grammar correctness* scores are very high (above 0.95) for all domains and methods, demonstrating that these neural models are capable of generating sentences free of grammatical mistakes. *Reading difficulty* estimates whether the generated utterances are more difficult to understand (2), have the same reading difficulty (1), or easier to understand (0) than the original utterance. That is, it ranges from 0 to 2, and lower values are better. We can see the scores are decreasing from the vanilla methods to ones with both type-guiding and reading difficulty filter. Thus, the reading difficulty filter indeed contributes achieving the task objectives. Finally, *type accuracy* measures whether the generated reformulation is consistent with the target reformulation type or not. We can see that the type-guided methods perform better in this regard.

Comparing the results with those from automatic evaluation, we find that the same observations hold regarding the improvements brought in by type-guiding and reading difficulty filtering, as well as generalizability (hybrid vs. single-domain performance).

### 6.3 Crowdsourcing Human Evaluation

To investigate the end-to-end performance of our model, we further evaluate the simulated utterances via crowdsourcing to test whether the generated sequence is indistinguishable from conversations performed by real users, following [53]. We sample 46 dialogues from the logged data, and replace the reformulation utterances in the original dialogues with our simulated utterances according to the same reformulation type. For example, in Table 7, the reformulations happens 2 times in the original dialogue (colored in blue and

**Table 7: An example reformulation dialogue piece and its corresponding simulated dialogue piece.**

Original Reformulation Piece	
User:	OK, now I want to know more about restaurants in Dubai
Agent:	Hotel? Apartment? or a bunk bed?
User:	Need to know about restaurants
Agent:	Sorry, I didn't get that. Can you rephrase?
User:	Restaurant
Simulated Reformulation Piece	
User:	OK, now I want to know more about restaurants in Dubai
Agent:	Hotel? Apartment? or a bunk bed?
User:	Can you find me a restaurant in Dubai?
Agent:	Sorry, I didn't get that. Can you rephrase?
User:	Places for dinner

red). We replace these original reformulations using our generated utterances. We randomly generate a total of 302 dialogue pieces. We place the original and simulated dialogue pieces side-by-side, in a random order, and present them to 3 crowd workers on Amazon MTurk. They are asked to choose which of the two dialogues is performed purely by a human; they are specifically instructed to compare the highlighted reformulated sequences. A tie is permitted when it is hard to distinguish. Additionally, workers are requested to give a brief explanation behind their choice. Options without explanations are filtered out. The results are presented in Table 5. We can see that the improvements of considering type-guided and reading difficulty filter are consistent with automatic and expert human evaluation. At the same time, the performance gaps are smaller. It is worth noting that 40% of pairs are not distinguishable, which also attests to the human-likeness of simulated responses.

### 6.4 Analysis

Next, we analyze the explanations given by crowd workers in Section 6.3 for their choice of which of the dialogues was performed purely by a human. We observe the same groups of reasons with [53] based on style and content.

For example, the aspects mentioned in terms of style include “this conversation seemed more natural and realistic to human conversation” (Realisticity), “the user keep trying to restart to get the right answer” (Engagement), and “this user seems to get a bit frustrated with the bot and adapts to it” (Emotion). Those based on content are include “the first one seems to understand the request a bit better” (Response), “the top is likely the human as there are punctuation and varied capitalizations” (Grammar), and “it says a lot of words that seem robotic” (Length).

There is also a new group of reasons focusing on the effect of reformulation. Most annotators take reformulation as being positive, e.g., “*they are asked to clarify their question, then they rephrase it so the Bot Agent can understand*.” But there are also negative cases, e.g., “*the user in 1 appears to get stuck in a loop and ends up requesting nonsensical topics*.” To summarize, integrating reformulation behavior contributes to making the simulation more human-like. However, low-quality reformulations, like sticking to nonsensical topics, might hurt.

## 7 CONCLUSIONS

This work has focused on one particular user trait, utterance reformulation in response to failed replies, to create more human-like user simulation. We have started our investigation with a user study, where interactions with five conversational agents across three domains were requested. The main outcome of the user study are seven common reformulation types that occur in conversational recommendations. To incorporate the findings from the user study into user simulation approaches, we have proposed a reformulation sequence generation task based on a given utterance. We have developed methods using type-guided text generation transformers to address this task, and have also included a way to estimate reading difficulty. The proposed methods have been compared using automatic evaluation, expert human evaluation, and crowdsourcing human evaluation. We have found that type-guiding or reading difficulty filtering alone does not always guarantee improvements, but their combination has been shown to improve performance significantly and substantially across methods, domains, and evaluation metrics. Furthermore, experiments on a hybrid-domain dataset have shown the generalizability of our method. The utterance reformulation dataset we have developed as part of this study consists of 1.8k reformulated utterances for 150 dialogues, which can be a useful resource for training more realistic simulators.

**Limitations.** Our user study has focused on modeling the reformulation behavior of a persistent user. In practice, a simulator would need to model users with different personality types, ranging from non-persistent to maximally persistent. Estimating the distribution of persistence for the simulated user population requires further investigation. We also note that our model is trained on a particular set of agents and might suffer from generalizability; for new agents, data may need to be collected for fine-tuning.

**Future work.** We see several avenues for extending our work in the future. First, the improved, more human-like reformulation behavior is to be incorporated into a larger simulation framework, like [53], and evaluated in the context of an end-to-end task. There are also additional dimensions of reformulation behavior to consider, such as the users’ knowledge and preferences, changing intents, style consistency of generated reformulations, and effects of the success rate of previous dialogue.

## ACKNOWLEDGMENTS

This work was partially supported by the National Natural Science Foundation of China (61973272).

## REFERENCES

- [1] Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W. Bruce Croft. 2019. Asking Clarifying Questions in Open-Domain Information-Seeking Conversations. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '19)*. 475–484.
- [2] Layla El Asri, Jing He, and Kaheer Suleman. 2016. A Sequence-to-Sequence Model for User Simulation in Spoken Dialogue Systems. In *Proceedings of the 17th Annual Conference of the International Speech Communication Association (Interspeech '16)*. 1151–1155.
- [3] Leif Azzopardi, Mateusz Dubiel, Martin Halvey, and Jeffery Dalton. 2018. Conceptualizing Agent-human Interactions during the Conversational Search Process. In *The Second International Workshop on Conversational Approaches to Information Retrieval (CAIR '18)*.
- [4] Krisztian Balog. 2021. Conversational AI from an Information Retrieval Perspective: Remaining Challenges and a Case for User Simulation. In *Proceedings of the 2nd International Conference on Design of Experimental Search & Information RETrieval Systems (DESIRE '21)*. 80–90.
- [5] Krisztian Balog, David Maxwell, Paul Thomas, and Shuo Zhang. 2021. Report on the 1st Simulation for Information Retrieval Workshop (Sim4IR 2021) at SIGIR 2021. *SIGIR Forum* 52, 2 (Dec. 2021), 11–26.
- [6] Krisztian Balog, Filip Radlinski, and Alexandros Karatzoglou. 2021. On Interpretation and Measurement of Soft Attributes for Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*. 890–899.
- [7] Anja Belz and Ehud Reiter. 2006. Comparing Automatic and Human Evaluation of NLG Systems. In *Proceedings of the 11th Conference of the European Chapter of the Association for Computational Linguistics (EACL '06)*. 313–320.
- [8] Morton B. Brown and Alan B. Forsythe. 1974. Robust Tests for the Equality of Variances. *J. Amer. Statist. Assoc.* 69, 346 (1974), 364–367.
- [9] Rocio Cañameres, Pablo Castells, and Alistair Moffat. 2020. Offline Evaluation Options for Recommender Systems. *Inf. Retr. J.* 23, 4 (2020), 387–410.
- [10] Kevyn Collins-Thompson and James P. Callan. 2004. A Language Modeling Approach to Predicting Reading Difficulty. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL '04)*. 193–200.
- [11] Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2020. TREC CAsT 2019: The Conversational Assistance Track Overview. arXiv:2003.13624 [cs.LR]
- [12] Jan Deriu, Álvaro Rodrigo, Arantxa Otegi, Guillermo Echegoyen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. 2019. Survey on Evaluation Methods for Dialogue Systems. arXiv:1905.04071 [cs.CL]
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (ACL '19)*. 4171–4186.
- [14] Chongming Gao, Wenqiang Lei, Xiangnan He, Maarten de Rijke, and Tat-Seng Chua. 2021. Advances and Challenges in Conversational Recommender Systems: A Survey. arXiv:2101.09459 [cs.CL]
- [15] Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural Approaches to Conversational AI. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts (ACL '18)*. 2–7.
- [16] David Griol, Javier Carbo, and José M. Molina. 2013. An Automatic Dialog Simulation Technique to Develop and Evaluate Interactive Conversational Agents. *Appl. Artif. Intell.* 27, 9 (oct 2013), 759–780.
- [17] Javeria Habib, Shuo Zhang, and Krisztian Balog. 2020. IAI MovieBot: A Conversational Movie Recommender System. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*. 3405–3408.
- [18] Jin Huang, Harrie Oosterhuis, Maarten de Rijke, and Herke van Hoof. 2020. Keeping Dataset Biases out of the Simulation: A Debaised Simulator for Reinforcement Learning Based Recommender Systems. In *Fourteenth ACM Conference on Recommender Systems (RecSys '20)*. 190–199.
- [19] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. 2020. A Survey on Conversational Recommender Systems. arXiv:2004.00646 [cs.LR]
- [20] Wenqiang Lei, Gangyi Zhang, Xiangnan He, Yisong Miao, Xiang Wang, Liang Chen, and Tat-Seng Chua. 2020. Interactive Path Reasoning on Graph for Conversational Recommendation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '20)*. 2073–2083.
- [21] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL '20)*. 7871–7880.
- [22] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Proceedings of the ACL workshop on Text Summarization Branches Out*. 74–81.
- [23] Aldo Lipani, Ben Carterette, and Emine Yilmaz. 2021. How Am I Doing?: Evaluating Conversational Search Systems Offline. *ACM Trans. Inf. Syst.* 39, 4, Article 51 (Aug. 2021), 22 pages.
- [24] Phillip Lippe, Pengjie Ren, Hinda Haned, Bart Voorn, and Maarten de Rijke. 2021. Diversifying Task-oriented Dialogue Response Generation with Prototype Guided Paraphrasing. *IEEE/ACM Trans. Audio Speech Lang. Process.* (2021). Submitted.

- [25] Qian Liu, Bei Chen, Jian-Guang Lou, Bin Zhou, and Dongmei Zhang. 2020. Incomplete Utterance Rewriting as Semantic Segmentation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP '20)*. 2846–2857.
- [26] Alexandros Papangelis, Yi-Chia Wang, Piero Molino, and Gokhan Tur. 2019. Collaborative Multi-Agent Dialogue Model Training Via Reinforcement Learning. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL '19)*. 92–102.
- [27] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: A Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL '02)*. 311–318.
- [28] Baolin Peng, Xiujun Li, Jianfeng Gao, Jingjing Liu, and Kam-Fai Wong. 2018. Deep Dyna-Q: Integrating Planning for Task-Completion Dialogue Policy Learning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL '18)*. 2182–2192.
- [29] Chen Qu, Liu Yang, W. Bruce Croft, Johanne R. Trippas, Yongfeng Zhang, and Minghui Qiu. 2018. Analyzing and Characterizing User Intent in Information-Seeking Conversations. In *Proceedings of the 41st International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '18)*. 989–992.
- [30] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. (2019).
- [31] Filip Radlinski, Krisztian Balog, Bill Byrne, and Karthik Krishnamoorthi. 2019. Coached Conversational Preference Elicitation: A Case Study in Understanding Movie Preferences. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL '19)*. 353–360.
- [32] Filip Radlinski and Nick Craswell. 2017. A Theoretical Framework for Conversational Search. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval (CHIIR '17)*. 117–126.
- [33] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. arXiv:1910.10683 [cs.LG]
- [34] Eun Rha, Wei Shi, and Nicholas Belkin. 2017. An Exploration of Reasons for Query Reformulations. *J. Assoc. Inf. Sci. Technol* 54 (2017), 337–346.
- [35] Corbin Rosset, Chenyan Xiong, Xia Song, Daniel Campos, Nick Craswell, Saurabh Tiwary, and Paul Bennett. 2020. Leading Conversational Search by Suggesting Useful Questions. In *Proceedings of The Web Conference 2020 (WWW '20)*. 1160–1170.
- [36] Alexandre Salle, Shervin Malmasi, Oleg Rokhlenko, and Eugene Agichtein. 2021. Studying the Effectiveness of Conversational Search Refinement Through User Simulation. In *Proceedings of the 43rd European Conference on Information Retrieval (ECIR '21)*. 587–602.
- [37] Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007. Agenda-based User Simulation for Bootstrapping a POMDP Dialogue System. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL '07)*. 149–152.
- [38] Ivan Sekulic, Mohammad Aliannejadi, and Fabio Crestani. 2022. Evaluating Mixed-initiative Conversational Search Systems via User Simulation. In *The 15th International Conference on Web Search and Data Mining (WSDM '22)*. 888–896.
- [39] Hui Su, Xiaoyu Shen, Rongzhi Zhang, Fei Sun, Pengwei Hu, Cheng Niu, and Jie Zhou. 2019. Improving Multi-turn Dialogue Modelling with Utterance ReWriter. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL '19)*. 22–31.
- [40] Weiwei Sun, Shuo Zhang, Krisztian Balog, Zhaochun Ren, Pengjie Ren, Zhumin Chen, and Maarten de Rijke. 2021. Simulating User Satisfaction for the Evaluation of Task-Oriented Dialogue Systems. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*. 2499–2506.
- [41] Yueming Sun and Yi Zhang. 2018. Conversational Recommender System. In *The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '18)*. 235–244.
- [42] Iuliana Teodorescu. 2009. Maximum Likelihood Estimation for Markov Chains. arXiv:0905.4131 [stat.CO]
- [43] Johanne R. Trippas, Damiano Spina, Lawrence Cavedon, Hideo Joho, and Mark Sanderson. 2018. Informing the Design of Spoken Conversational Search: Perspective Paper. In *Proceedings of the 2018 Conference on Human Information Interaction and Retrieval (CHIIR '18)*. 32–41.
- [44] Johanne R. Trippas, Damiano Spina, Falk Scholer, Ahmed Hassan Awadallah, Peter Bailey, Paul N. Bennett, Ryan W. White, Jonathan Lio, Yongli Ren, Flora D. Salim, and Mark Sanderson. 2019. Learning About Work Tasks to Inform Intelligent Assistant Design. In *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval (CHIIR '19)*. 5–14.
- [45] Svitlana Vakulenko, Shayne Longpre, Zhucheng Tu, and Raviteja Anantha. 2021. Question Rewriting for Conversational Question Answering. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM '21)*. 355–363.
- [46] Alexandra Vtyurina, Charles L. A. Clarke, Edith Law, Johanne R. Trippas, and Horatiu Bota. 2020. A Mixed-Method Analysis of Text and Audio Search Interfaces with Varying Task Complexity. In *Proceedings of the 2020 ACM SIGIR on International Conference on Theory of Information Retrieval (ICTIR '20)*. 61–68.
- [47] Alexandra Vtyurina and Adam Fournay. 2018. Exploring the Role of Conversational Cues in Guided Task Support with Virtual Assistants. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. 1–7.
- [48] Alexandra Vtyurina, Denis Savenkov, Eugene Agichtein, and Charles L. A. Clarke. 2017. Exploring Conversational Search With Humans, Assistants, and Wizards. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. 2187–2193.
- [49] Xiaojun Wan. 2018. Automatic Text Simplification by Horacio Saggion. *J. Comput. Linguist* 44, 4 (2018), 659–661.
- [50] Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2018. Neural Network Acceptability Judgments. arXiv:1805.12471 [cs.CL]
- [51] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations (EMNLP '20)*. 38–45.
- [52] Hamed Zamani, Johanne R. Trippas, Jeff Dalton, and Filip Radlinski. 2022. Conversational Information Seeking. arXiv:2201.08808 [cs.IR]
- [53] Shuo Zhang and Krisztian Balog. 2020. Evaluating Conversational Recommender Systems via User Simulation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '20)*. 1512–1520.