

# Limitations of Current Evaluation Practices for Conversational Recommender Systems and the Potential of User Simulation

Nolwenn Bernard\*  
TH Köln  
Cologne, Germany  
nolwenn.bernard@th-koeln.de

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

## Abstract

Research and development on conversational recommender systems (CRSs) critically depends on sound and reliable evaluation methodologies. However, the interactive nature of these systems poses significant challenges for automatic evaluation. This paper critically examines current evaluation practices and identifies two key limitations: the over-reliance on static test collections and the inadequacy of existing evaluation metrics. To substantiate this critique, we analyze real user interactions with nine existing CRSs and demonstrate a striking disconnect between self-reported user satisfaction and performance scores reported in prior literature. To address these limitations, this work explores the potential of user simulation to generate dynamic interaction data, offering a departure from static datasets. Furthermore, we propose novel evaluation metrics, based on a general reward/cost framework, designed to better align with real user satisfaction. Our analysis of different simulation approaches provides valuable insights into their effectiveness and reveals promising initial results, showing improved correlation with system rankings compared to human evaluation. While these findings indicate a significant step forward in CRS evaluation, we also identify areas for future research and refinement in both simulation techniques and evaluation metrics.

## CCS Concepts

• Information systems → Recommender systems.

## Keywords

Conversational recommender systems; Evaluation; User simulation

## ACM Reference Format:

Nolwenn Bernard and Krisztian Balog. 2025. Limitations of Current Evaluation Practices for Conversational Recommender Systems and the Potential of User Simulation. In *Proceedings of the 2025 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region (SIGIR-AP 2025)*, December 7–10, 2025, Xi'an, China. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3767695.3769478>

\*Work done while at the University of Stavanger, Norway



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGIR-AP 2025, Xi'an, China*

© 2025 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2218-9/2025/12  
<https://doi.org/10.1145/3767695.3769478>

## 1 Introduction

The evaluation of conversational recommender systems (CRSs) is a complex problem, requiring consideration of both system effectiveness and user experience. There is a long-standing tradition in recommender systems of adapting evaluation methodologies from information retrieval (IR), a trend seen, for instance, in the evolution from rating prediction to top-k ranking evaluation [6, 13, 37, 43]. However, for CRSs, this established paradigm proves insufficient, as returning relevant recommendations at top ranks does not always guarantee user satisfaction [34]. The interactive setting further contributes to the complexity of the evaluation process. While different evaluation paradigms have been proposed, a standard methodology remains to be established [16, 21].

In this work, we investigate current practices for automatically evaluating the effectiveness of CRSs, typically performed offline using benchmark conversational datasets (e.g., ReDial [26] and OpenDialKG [32]). We identify two main limitations of this approach: (1) reliance on static datasets that overlook the interactive nature of conversations and (2) the inadequacy of traditional metrics to assess user *utility*. To address these fundamental limitations and contribute to establishing a more appropriate evaluation paradigm, our primary contributions are twofold: first, we introduce novel, user-centric metrics grounded in the economic principles of reward and cost during interaction, moving beyond traditional effectiveness measures. Second, we perform an empirical investigation into the use of user simulation for generating synthetic conversations, thereby enabling a more dynamic and automated evaluation framework for CRSs.

Our first contribution, focusing on developing more appropriate user utility metrics, stems from the observation that current evaluation metrics for CRSs often prioritize specific system attributes like recommendation accuracy and fluency, neglecting the ultimate goal of user utility. Moreover, these metrics are frequently borrowed from static settings (e.g., Recall@N from IR, Distinct-n from natural language generation) and are therefore ill-suited for capturing the interactive nature of CRSs. Indeed, we empirically demonstrate that current CRS evaluation based on static test collections and traditional metrics fail to assess the actual utility generated for users. We advocate for a user-centric evaluation approach, emphasizing utility. Consequently, we propose a utility framework, inspired by economic models of human-computer interaction [3, 4], that balances the reward gained and cost incurred for a user interacting with a CRS. Based on this framework, we propose two new metrics and demonstrate their stronger alignment with user satisfaction compared to traditional metrics like Recall@N. Furthermore, we re-examine previously reported CRS performance

results [29, 46] using our best-performing utility-based metric and find that some conclusions drawn in prior work no longer hold.

Regarding the reliance on static datasets that overlook the interactive nature of conversations, user simulation offers a potential path towards automated CRS evaluation. Its widespread adoption, however, is hindered by questions surrounding the validity of user simulators [5]. Nevertheless, previous work has shown that user simulators can reproduce the ranking of CRSs obtained from offline experiments and user studies [20, 46, 52], supporting the viability of simulation-based evaluation as an alternative to traditional offline experiments. Building on this promising avenue, our second key contribution involves a comparative analysis of different user simulation approaches. To the best of our knowledge, a direct comparison of user simulators for evaluating CRSs and their alignment with human judgments has not yet been performed. Following the framework proposed in [8], we compare the reliability of traditional agenda-based and more recent LLM-based simulators in evaluating nine existing CRSs against CRS rankings in terms of self-reported user satisfaction. Our findings reveal that the LLM-based simulator is more reliable; however, limitations remain in terms of reproducing the relative ranking of a diverse set of CRSs.

To summarize, this paper makes four key contributions that address the limitations of current CRS evaluation. First, we introduce a general reward-cost framework for measuring user utility in CRS interactions. Second, we propose two novel metrics derived from this framework and demonstrate their positive correlation with self-reported user satisfaction. Third, by re-evaluating the performance of existing CRSs using our best-performing utility metric, we show that certain findings from prior work based on static datasets are no longer supported. Finally, we present a comparative analysis of different user simulator approaches regarding their reliability in evaluating CRSs. Overall, this work takes a critical step towards more ecologically valid and user-centric automated CRS evaluation utilizing user simulation.

## 2 Related Work

This work focuses on the evaluation of conversational recommender systems (CRSs) and the potential of user simulation to alleviate some of the limitations of current offline evaluation practices. In this section, we provide an overview of CRSs types and evaluation methodologies, and user simulation approaches.

### 2.1 Conversational Recommender Systems

We identify four main types of CRS proposed over the years. First, *rule-based CRSs* rely on a set of predefined rules to generate responses [14, 17, 42]. This allows for total control over the system’s behavior as well as transparency. However, it also presents disadvantages such as the difficulty of maintaining and updating the rules as the system extends to new domains. Second, *retrieval-based CRSs* follow the idea of retrieving the most relevant response from a corpus of predefined responses, including placeholders for recommended items [30]. This approach is more flexible than the rule-based one, while keeping high quality responses with regards to naturalness and meaningfulness. However, having predefined responses limits the diversity and adaptability of the system. Third, CRSs can be built using deep neural networks, following either a

pipeline or end-to-end architecture [11, 47, 55]. These systems are more flexible to unknown contexts and easier to adapt to new domains, given sufficient training data. However, using deep neural networks reduces the transparency of the system and its controllability due to the black-box nature of these models. Fourth, the most recent type of CRS relies on large language models (LLMs) [15, 19, 38]. LLMs have shown great potential in generating human-like responses which is critical in a conversational setting. The recommendation task may be performed by the LLM itself based on its internal knowledge or in combination with a recommender system. LLM-based CRSs require less even no data for training considering the pre-trained models available. Hence, they appear as data-efficient and accessible solutions for conversational recommendation. However, similarly to neural-based CRSs, their transparency and controllability is limited.

### 2.2 Evaluation Methodologies

Traditionally, the main methodologies to evaluate CRSs are: online, user studies, and offline evaluation [5]. *Online evaluation* requires a real-world setting where real users interact with a CRS, thus, providing arguably the most reliable evaluation. In practice, it often corresponds to A/B tests where the performances of different versions of a CRS are compared with regards to a set of performance metrics. Such evaluation can be difficult to conduct, especially in an academic setting, because it requires access to a sufficient number of real users and different versions of the CRS to compare. *User studies* involve human participants that are not necessarily actual users of the system. These studies, conducted in a laboratory or crowdsourcing setting, allow for the assessment of a CRS with regards to different characteristics at different granularity levels [21]. For example, Manzoor et al. [31] and Manzoor and Jannach [29] study the meaningfulness of system responses, while, Li et al. [26] assess the quality of system response at the conversation level. Finally, *offline evaluation* does not rely on human participants and is commonly used to evaluate algorithms and models using automatic metrics such as Recall and BLEU. Recently, LLMs are being employed as an alternative to human judges for text evaluation [12, 18]. This makes it possible to automatically evaluate more “subjective” aspects, such as the quality of explanations [20] and conciseness [18].

*Simulation-based evaluation*, which is gaining attention in the field, bridges the gap between online/user studies and offline evaluation. Indeed, user simulators aim to mimic the behavior of human users, therefore, simulated users can substitute real users to perform the evaluation of CRSs in a controlled and interactive manner [5]. There are two main strategies to using simulation: (1) the user simulator interacts with the CRS and provides direct feedback on the system’s performances [39], and (2) the user simulator interacts with the CRS to generate dialogues that are separately evaluated [20, 46, 52]. Despite recent advances, simulation-based evaluation is still dependent on benchmark datasets. For example, Wang et al. [46] generate synthetic dialogues based on excerpts of dialogues from ReDial [26] or OpenDialKG [32] and stop either when a ground truth item is recommended or a predetermined number of turns is reached, leading to potentially unnatural endings.

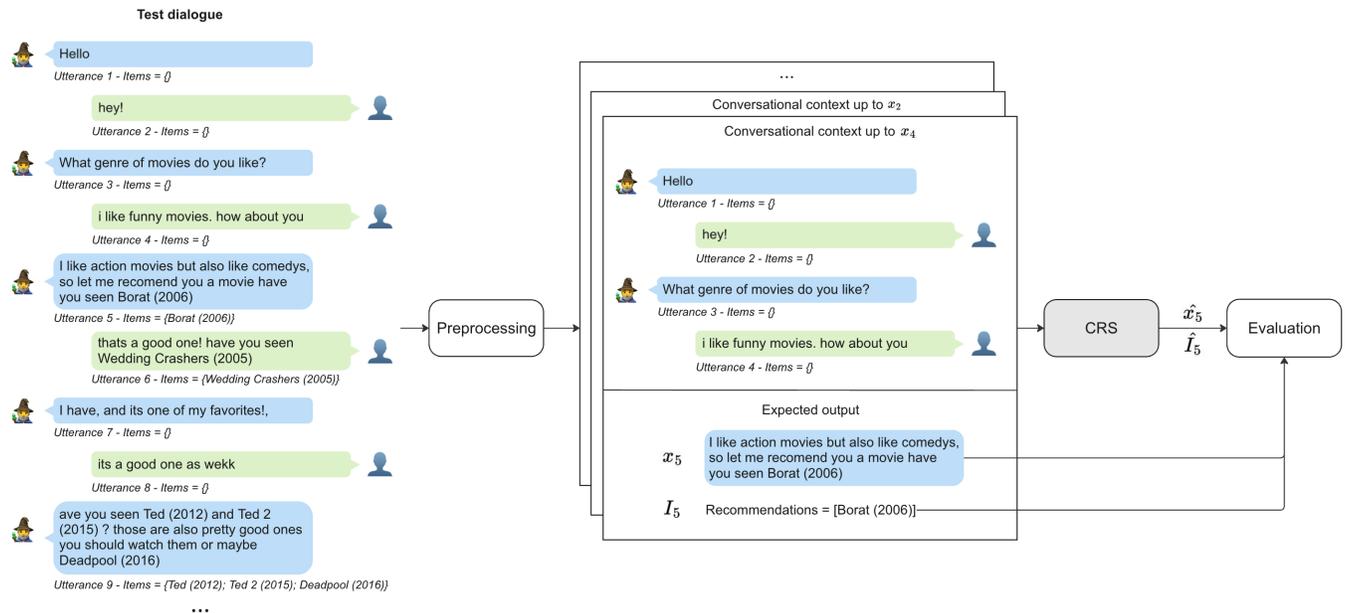


Figure 1: Illustration of the current CRS evaluation practice, based on static conversational datasets.

### 2.3 User Simulation Approaches

User simulators share architectural similarities with CRSs, given that both are intelligent agents that need to generate utterances, which involves the tasks of natural language understanding, dialogue management, and natural language generation. We identify three types of user simulator: agenda-based, deep neural-based, and LLM-based. *Agenda-based user simulators* determine the next action based on an agenda built from a goal and that is updated based on the context of the conversation [35]. A Markovian state representation is assumed, and the allowed state transitions are specified by an interaction model. This design allows for transparency of the simulated users’ behavior. However, the creation of such user simulators requires a certain knowledge of the domain and task, especially for the creation of the interaction model. To address this issue, *data-driven user simulators* have been proposed, which often rely on deep neural networks to learn user behavior from data directly [28]. This type of user simulators often abstracts the response generation process, making it less interpretable. Similarly to neural-based CRSs, the transparency of these simulators is limited and they require sufficient data for training. Finally, the rise of LLMs has also impacted the creation of user simulators. Indeed, recent works have proposed to build user simulators based on LLMs for conversational information access and task-oriented dialogue systems [1, 20, 27, 36, 46, 50]. For example, Abbasiantaeb et al. [1] use LLM-based user simulators for data generation in the context of conversational question answering, while Huang et al. [20] and Wang et al. [46] use LLM-based user simulators for CRS evaluation.

## 3 Current Evaluation Practices and their Limitations

This section first details the predominant approach to automatic CRS evaluation, and then identifies two key limitations: the reliance on offline datasets and the choice of metrics.

### 3.1 Problem Statement

In this work, we focus on the evaluation of CRSs. The objective of a CRS is to assist users in satisfying recommendation-related goals through a multi-turn dialogue [22]. The CRS and the user are active dialogue participants that take turns in issuing utterances  $x_i$ . Therefore, a dialogue is defined as a sequence of utterances:  $d = \{x_1, x_2, \dots, x_n\}$ . We note that the utterances can be annotated with semantic information, such as dialogue acts [23], as well as metadata. The metadata typically includes items  $I$  that are recommended by the CRS or mentioned by the user.

### 3.2 Current Practice

Offline evaluation of CRSs typically involves three stages, leveraging annotated conversational datasets, illustrated in Fig. 1. First, dialogues from the test portion of the dataset (hereinafter: test dialogues) are preprocessed and/or augmented to derive inputs and expected outputs for the CRS. The input corresponds to the conversational context up to utterance  $x_t$  (issued by the user) along with optional metadata (e.g., entity mentioned). The expected output corresponds to the next utterance  $x_{t+1}$  and associated recommendations  $I_{t+1}$  (which may be empty) in the test dialogue. Second, the CRS generates the next utterance  $\hat{x}_{t+1}$  and the associated recommendations  $\hat{I}_{t+1}$  given the input. Third, evaluation metrics are computed on the *turn level*, and then aggregated (micro-averaged or macro-averaged). Traditionally, the metrics assess either *task*

*performance* in terms of recommendation accuracy, by comparing  $\hat{I}_{t+1}$  with the ground truth target items  $I_{t+1}$ , or the quality of the generated utterance  $\hat{x}_{t+1}$  either with respect to the expected output utterance  $x_{t+1}$  or on its own account (e.g., fluency). This process is repeated and the results are aggregated over all test dialogues.

### 3.3 Limitation #1: Reliance on Offline Datasets

The previous description of the traditional offline evaluation process highlights the static nature of the datasets used. By definition, they neglect the interactive nature of conversational systems [5]. The recommended items in  $I_{t+1}$ , which are often annotated after the data collection, are considered as ground truth. We observe that typically,  $I_{t+1}$  comprises a single item, for example, it represents more than 80% of the cases in the ReDial dataset [26]. This is unrealistic, as a user is unlikely to look for a specific item when interacting with a CRS, but rather for recommendations that satisfy their current preferences, consequently the effectiveness of the CRSs may be underestimated [44]. Moreover, He et al. [19] notice that items in  $I$  may be repeated during the conversation, i.e., ground truth items comprise items that were already mentioned in the conversation, which can lead to overestimating the effectiveness of CRSs. Furthermore, the lack of diversity in dialogue patterns within offline datasets has also been noted [21]. To mitigate these limitations, in this work, we use dialogues collected with real users and synthetic dialogues generated with user simulators that are guided by an overall goal [35] rather than specific target item(s).

### 3.4 Limitation #2: Evaluation Metrics

Automatic metrics used in offline evaluation often focus on specific system aspects, such as recommendation accuracy or linguistic diversity, rather than overall user utility. For instance, Recall@N is a commonly used metric to evaluate recommendation accuracy (see, e.g., [11, 20, 46, 47, 55]), typically computed with  $N$  set to 1, 10, 25, and 50. Our review of the literature does not provide a clear rationale for the choice of large  $N$  values ( $N > 10$ ). We argue that in a conversational setting, where utterances should remain focused and concise, users expect targeted and manageable suggestions rather than a large batch of items that are likely to overwhelm them. However, to the best of our knowledge, this point has not been previously discussed in the literature, and these values for  $N$  have been widely adopted in the field. We hypothesize that the prevalent use of larger  $N$  values might stem from the common assumption of a single relevant ground truth item in the evaluation dataset (as discussed in Limitation #1). In such scenarios, lower  $N$  values could indeed lead to metrics that are too insensitive to effectively differentiate between recommendation algorithms. However, this methodological choice carries the significant risk of the evaluation metric becoming increasingly detached from real-world user experience and satisfaction. Indeed, in recent work, Bernard et al. [9] found a negligible correlation between Recall@10 and self-reported user satisfaction from real users, suggesting that Recall@N may not be a reliable proxy for user utility. However, they do not propose alternative metrics. Similarly, some studies report distinct-n to assess the linguistic diversity of the system without considering the conversational context [11, 47, 55]. However, this

approach, in our view, does not accurately reflect a CRS’s ability to assist users in finding relevant items. Furthermore, some of the commonly used metrics were originally designed for static tasks, such as recommendation or language generation, and thus often overlook the interactive nature of conversational systems. In this work, we propose alternative metrics that evaluate CRSs with regards to user utility, considering the conversational nature of the task and lifting the assumption of the user looking for a single specific target items.

## 4 Evaluation Metrics

The ultimate goal of a conversational recommender system is to generate utility for its users [21], which should ideally manifest as high user satisfaction. However, as demonstrated in [9], and also confirmed by our experiments in Section 5, the widely used Recall@N measure fails to adequately capture this utility. This leads to our central research question for this section: *Can we develop evaluation metrics for CRSs focused on user utility that show a greater correlation with overall end-user satisfaction?*

To address this question, we first propose a formal framework to provide a more direct measure of user utility by considering both the effort expended by the user (i.e., cost) during their interactions with the CRS and the benefit (i.e., reward) gained from the recommendations made by the system (Section 4.1). We show how this framework can encapsulate existing evaluation measures (Section 4.2) and then propose novel evaluation metrics that aim to provide better approximations of costs and rewards (Section 4.3).

### 4.1 A General Reward/Cost Framework

Evaluating the effectiveness of a CRS requires considering not only the benefits users derive but also the costs associated with the interaction. To this end, we propose a general utility framework for measuring the effectiveness of a CRS inspired by economic models of human-computer interaction [3, 4]. Our measure of utility is based on the balance between the *reward*  $R$  the user is gaining from interacting with the system and the overall *cost*  $C$  incurred throughout the conversation.

We assume that there are multiple reward factors  $r_i$ , each with its associated weight  $\alpha_i$ , that can contribute to the overall reward:  $R = \sum_i \alpha_i r_i$ . For example, the CRS making relevant recommendations, exposing the user to diverse items, or enabling serendipitous discoveries are all positive outcomes where the user gains from interacting with the system. Similarly, we write cost as a weighted linear combination of various cost factors:  $C = \sum_j \beta_j c_j$ . Examples of cost factors include the interaction length (total number of turns), the cognitive load associated with understanding the system’s questions and formulating responses, and the time the user has to wait for the system to respond.

The utility of the conversational recommendation process is then defined as the ratio of the total reward to the total cost:

$$\text{utility} = \frac{R}{C}. \quad (1)$$

A higher utility score indicates a more effective CRS, where the user gains significant reward with relatively low interaction costs.

Utility, as defined in Eq. (1) for a single conversation, is then aggregated into a final system-level metric by averaging its value across a set of evaluation conversations.

We note that this framework offers a general approach applicable to various interactive systems. However, our primary focus in this work lies specifically in its instantiation for the task of conversational recommendations. In the subsequent sections, we will discuss relatively simple instantiations of this framework, considering a single reward factor and a single cost factor at a time. This simplification allows us to demonstrate the potential of measuring effectiveness from an economic perspective, without the need to define specific weight coefficients ( $\alpha_i$  and  $\beta_j$ ). The development and analysis of more advanced metrics incorporating multiple weighted factors and normalization techniques are left for future work.

## 4.2 Existing Measures

We show how two of the most commonly used CRS effectiveness measures can be instantiated in our reward/cost framework.

**Recall@N.** This measure is only computed for turns where a recommendation has been made. Then, the turn-level recall values are aggregated on the dataset level,<sup>1</sup> corresponding to micro-averaging. As such, it does not directly measure the effectiveness of in a given dialogue, which explains why it does not correlate with user satisfaction. All other measures we discuss in this paper perform macro-averaging instead, by first evaluating conversations and then averaging over all conversations in the test dataset.

In terms of the Reward/Cost framework, the total cost, across all turns  $t$  in all test dialogues  $d$ , is the total number of item recommendations made. The total reward is the number of recommended items that match the ground truth target items. Formally:

$$\text{Recall@N} = \frac{1}{\sum_d |d|} \frac{\sum_d \sum_t |I_t \cap \hat{I}_t|}{\sum_d \sum_t N \mathbb{1}_{rec}(t)}, \quad (2)$$

where  $|d|$  is the number of turns in a dialogue and  $\mathbb{1}_{rec}(t)$  is an indicator function that is 1 if a recommendation was made in turn  $t$  and 0 otherwise.

**Success Rate (SR).** A simple and intuitive way to assess a CRS is based on the ratio of “successful” dialogues in the test dataset  $D$ . Success is usually determined by checking if any recommendation matches any of the ground truth target items [24, 40, 49, 51]. That is, the reward for a single conversation is a binary value indicating whether there was a successful recommendation or not:  $R = \mathbb{1}_{succ}(d)$ . The cost is not considered, i.e., it’s a constant for all dialogues:  $C = 1$ . Putting these choices together, success rate on a set of dialogues  $D$  is measured as:

$$SR = \frac{\sum_d \mathbb{1}_{succ}(d)}{|D|}. \quad (3)$$

<sup>1</sup>This is not explicitly stated in any of the papers reporting on Recall@N and is based on our understanding of the implementation in CRSLab [54], a general toolkit to develop and evaluate CRSs (see, e.g., [47, 56]).

## 4.3 Novel Measures

Current CRS evaluation is limited by its dependence on ground truth target items (typically a single item). Additionally, these metrics largely ignore the cost incurred by the user during the interaction, or only consider it selectively during recommendation turns. We propose two novel user-centric metrics that consider the underlying intent of the user utterances. We operationalize this intent by leveraging the concept of a *dialogue act*, which is a structured representation of the of the communicative goal behind an utterance [23].

**Successful Recommendation Round Ratio (SRRR).** Success rate is a binary metric that does not account for multiple recommendation rounds in a dialogue, which can impact the cognitive load for user (i.e., more recommendations to assess). We introduce the notion of a *recommendation round* in a dialogue: a *round* starts when the intent of the CRS is Recommend and finishes when the user makes a decision to either Accept or Reject the recommendation. Each recommendation round has a fixed cost  $C = |\text{rounds}|$ . Successful recommendation rounds (ending with an acceptance) yield a reward of 1, otherwise 0:  $R = \mathbb{1}_{acc}(\text{round})$ . Putting these choices together:

$$SRRR = \frac{1}{|D|} \sum_d \frac{\sum_{\text{round}} \mathbb{1}_{acc}(\text{round})}{|\text{rounds}|}. \quad (4)$$

**Reward-per-Dialogue-Length (RDL).** Notice that all the previous measures either do not consider cost at all or only consider the turns where a recommendation has been made. Instead, we argue to consider all conversation turns as part of the effort expended by the user. A simple way of doing that is setting the cost to the total number of turns in the dialogue:  $C = |d|$ . The reward is the total number of recommendations accepted by the user  $R = |I_{acc}|$ . Then, the reward-per-dialogue-length on a set of dialogues is:

$$RDL = \frac{1}{|D|} \sum_d \frac{|I_{acc}|}{|d|}. \quad (5)$$

## 5 Experimental Evaluation

This section presents an experimental comparison of various evaluation metrics, focusing on their correlation with self-reported user satisfaction. Furthermore, we examine the performance of the different CRSs when interacting with real users.

### 5.1 Experimental Setup

For our experiments, we utilize conversations from the CRSArena-Dial dataset [9]. This dataset comprises 467 conversations real users conducted with nine different CRSs; see Table 1 for a breakdown per CRS. Additionally, this dataset contains self-reported user satisfaction (i.e., satisfied or frustrated), and direct and free-form feedback from users with regards to their experience with the CRSs.

In order to compute the different utility metrics, we automatically annotate the utterances with dialogue acts. For example, the utterance “Have you seen Step Brothers or Walk Hard: The Dewey Cox Story” can be annotated with the following dialogue act: `Recommend(TITLE=`Step Brothers`, TITLE=`Walk Hard: The Dewey Cox Story`)`. The annotation is performed following a few-shot

**Table 1: Summary of dialogue corpora in CRSArena-Dial [9].**

CRS	#dialogues
CRB-CRS_ReDial	60
KBRD_OpenDialKG	59
KBRD_ReDial	61
BARCOR_OpenDialKG	55
BARCOR_ReDial	46
UniCRS_OpenDialKG	42
UniCRS_ReDial	48
ChatCRS_OpenDialKG	44
ChatCRS_ReDial	52
Total	467

prompting approach with the Mistral NeMo<sup>2</sup> large language model; the intent schema is taken from the IARD dataset [10].

Given that we do not have ground truth item annotations, we take Recall@1 values from [45] and measure whether a recommendation was successful or not ( $R = \mathbb{1}_{succ}(d)$  in Eq. 3) based on the presence of dialogue acts with the intent Accept in user utterances.

## 5.2 Conversational Recommender Systems

We consider nine CRSs based on five different architectures, tailored for the movie recommendation domain.<sup>3</sup> The CRSs are trained and/or leverage external knowledge from either a subset of OpenDialKG [32] or ReDial [26]. The selection of these CRSs is guided by their public availability and use in previous studies. While we acknowledge limitations such as the emphasis on publicly available English-speaking CRSs, this selection represents a broad range of systems based on existing research.

- **CRB-CRS** [30] comprises two main components: (1) a retrieval and ranking component and (2) a recommendation and metadata integration component. The first is responsible for finding the best pre-defined response candidate from a pool based on the last user utterance and conversation history. The second adapts the previously retrieved response candidate by replacing placeholders with the recommended items and other relevant information.
- **KBRD** [11] uses knowledge propagation to integrate recommendation and dialogue systems for mutual enhancement. The recommender system takes advantage of the dialogue context to provide better recommendations, and the dialogue system uses the information about the recommended items and user model to generate more consistent responses. KBRD uses an external knowledge graph to enhance the representation of entities in the dialogue context and recommended items. The propagation of knowledge-enhanced representations helps to find connected items to recommend and to generate more informative responses.
- **BARCOR** [45] performs the recommendation and dialogue generation tasks using a single model, based on the pre-trained language model BART [25]. Specifically, it uses the encoder for the recommendation task and the decoder to generate responses. BART

is fine-tuned with examples where a specific token is used to mask the recommended items.

- **UniCRS** [47] uses the prompt learning paradigm to fulfill the tasks of recommendation and conversation in a unified manner. UniCRS relies on knowledge-enhanced prompts which include fused knowledge representation, task specific tokens, and dialogue context, in addition to a response template for the recommendation task. These prompts are then processed by the pre-trained language model DialoGPT [53].
- **ChatCRS** [46] combines ChatGPT and a recommendation model to perform conversational recommendation. The recommendation model computes the similarity between the embedded conversational history and embedded candidate items. Item embeddings are created using the model *text-embedding-ada-002* [33] from item attributes such as title, genre, and actor for movies. While ChatCRS can directly recommend items, the use of candidate items limits hallucinations and recommendation of items outside of the evaluation dataset.

## 5.3 Results

We evaluate the different utility metrics based on their correlation with self-reported user satisfaction. That is, we compute the Kendall’s  $\tau$  correlation between the rankings of CRSs obtained with each evaluation metric presented in Section 4 and the ranking based on self-reported user satisfaction. The results, reported in Table 2, show that *Success Rate* (SR), *Successful Recommendation Round Ratio* (SRRR), and *Reward-per-Dialogue-Length* (RDL) are positively correlated with self-reported user satisfaction, with Kendall’s  $\tau$  values of 0.67, 0.32, and 0.78, respectively. It suggests that, in our case, Reward-per-Dialogue-Length is the best proxy for user satisfaction among the utility metrics considered. Interestingly, we observe that Success Rate has a higher correlation with user satisfaction than Successful Recommendation Round Ratio, despite the latter being a more fine-grained measure. This may be explained by that SRRR only considers recommendation rounds as part of the cost and the identification of recommendation rounds might be prone to errors due to imperfections of automatic dialogue act annotation.

## 5.4 Analysis

Having established Reward-per-Dialogue-Length (RDL) as the best proxy for user satisfaction, we further analyze the relative rankings of the CRSs based on self-reported user satisfaction and RDL, and evaluate the validity of prior claims made using Recall@N. We make the following observations:

- We observe that a CRS with the same architecture trained with ReDial is generally better than the one trained with OpenDialKG according to both user satisfaction and RDL, with the exception of *ChatCRS*. It suggests that ReDial is a more suitable dataset for training a CRS in the movie recommendation domain. This contrasts with the ranking based on Recall@1 where the CRSs trained with OpenDialKG outperform those trained with ReDial.
- User satisfaction is generally low, with a maximal satisfaction of 0.523 for *ChatCRS\_OpenDialKG*, which is consistent with the low RDL scores observed. This shows a clear inability of the CRSs to

<sup>2</sup><https://mistral.ai/news/mistral-nemo/>

<sup>3</sup>We use their implementation from CRSArena at: <https://huggingface.co/spaces/iai-group/CRSArena/tree/main>.

**Table 2: Evaluation of CRSs with respect to self-reported user satisfaction (taken from [9]), which is taken as the ground truth (top row). Correlations are measured using Kendall’s  $\tau$ . The retrieval-based CRS is highlighted in bold blue and the CRSs with the best architecture according to previous work [20, 31, 46] are highlighted in bold orange.**

Metric	Kendall’s $\tau$	Ranking
User satisfaction	–	<b>ChatCRS_OpenDialKG</b> (0.523) > <b>ChatCRS_ReDial</b> (0.453) > BARCOR_ReDial (0.298) > BARCOR_OpenDialKG (0.145) > UniCRS_ReDial (0.102) > KBRD_ReDial (0.081) > <b>CRB-CRS_ReDial</b> (0.079) > UniCRS_OpenDialKG (0.048) > KBRD_OpenDialKG (0.017)
Recall@1 <sup>†</sup>	0.07	BARCOR_OpenDialKG (0.312) > <b>ChatCRS_OpenDialKG</b> (0.310) > UniCRS_OpenDialKG (0.308) > KBRD_OpenDialKG (0.231) > UniCRS_ReDial (0.050) > <b>ChatCRS_ReDial</b> (0.037) > BARCOR_ReDial (0.031) > KBRD_ReDial (0.028)
Success rate	0.67	<b>ChatCRS_OpenDialKG</b> (0.114) > <b>CRB-CRS_ReDial</b> (0.111) > <b>ChatCRS_ReDial</b> (0.094) > BARCOR_OpenDialKG (0.091) > BARCOR_ReDial (0.085) > UniCRS_ReDial (0.082) > KBRD_ReDial (0.048) > UniCRS_OpenDialKG (0.024) > KBRD_OpenDialKG (0.017)
Successful rec. round ratio	0.32	<b>CRB-CRS_ReDial</b> (0.077) > <b>ChatCRS_ReDial</b> (0.031) > BARCOR_OpenDialKG (0.030) > BARCOR_ReDial (0.023) > UniCRS_ReDial (0.021) > <b>ChatCRS_OpenDialKG</b> (0.018) > KBRD_OpenDialKG (0) > KBRD_ReDial (0) > UniCRS_OpenDialKG (0)
Reward-per-Dialogue-Length	0.78	<b>ChatCRS_OpenDialKG</b> (0.037) > <b>ChatCRS_ReDial</b> (0.025) > BARCOR_ReDial (0.024) > <b>CRB-CRS_ReDial</b> (0.023) > UniCRS_ReDial (0.019) > BARCOR_OpenDialKG (0.018) > KBRD_ReDial (0.008) > UniCRS_OpenDialKG (0.006) > KBRD_OpenDialKG (0.004)

<sup>†</sup> CRB-CRS\_ReDial is not included as it is not evaluated in [46].

efficiently assist users in fulfilling their needs and highlights the need for future work.

- The results show that *ChatCRS* outperforms the other CRSs in terms of user satisfaction and RDL. This is consistent with conclusions from previous work [20, 31, 46].
- We observe that *CRB-CRS\_ReDial*, which uses retrieval-based responses, is in the top half of the rankings for RDL. In other words, it outperforms some neural CRSs using generation-based responses, such as *UniCRS\_ReDial* and *KBRD\_ReDial*. However, *CRB-CRS\_ReDial* is in the lowest half for self-reported user satisfaction. Moreover, in this ranking, *KBRD\_ReDial* is better than *CRB-CRS\_ReDial*, it contrasts with the findings from the user study in [29], which showed that *CRB-CRS* outperforms *KBRD* in terms of response’s meaningfulness. Direct feedback from *CRSArena-Dial* provides some insights on the difference between the rank regarding RDL and user satisfaction, that is, some users reported that *CRB-CRS\_ReDial* was able to provide some relevant recommendations, but other factors such as the response time, stubbornness, and misunderstandings led to a feeling of frustration.

## 6 User Simulation-based Evaluation

A significant bottleneck in the evaluation of CRSs lies in the reliance on human users to generate interaction data. Collecting and analyzing conversations with real users is a time-consuming and resource-intensive process, hindering the scalability of evaluation efforts, particularly when comparing numerous systems. This motivates our research question: *Can user simulation offer a viable solution to automate the evaluation process, and if so, to what extent can simulations reliably predict the performance of CRSs with real users?* This section delves into the potential of user simulation as a means of automating the collection of interaction data for CRS evaluation, exploring the capabilities and limitations of two fundamentally different simulation approaches in mirroring real user conversations.

### 6.1 User Simulators

We consider two user simulators for movie recommendation:

- **ABUS** is an agenda-based user simulator with a traditional modular architecture comprising modules for natural language understanding and generation and a dialogue manager. In this case, natural language understanding and generation are performed with Mistral NeMo using a few-shot prompting approach, where the examples are taken from the IARD dataset [10]. To create and update the agenda, the dialogue manager relies on an interaction model built using a combination of handcrafted and data-driven rules. The data-driven rules are derived from a subset of the IARD dataset [10].
- **LLM-US** is a simulator relying on the Mistral NeMo LLM. In practice, upon reception of an utterance, *LLM-US* first decides to either continue the conversation or abort it because it does not progress towards the user goal. If the conversation is not aborted, *LLM-US* generates the next utterance using a zero-shot prompting approach using a prompt inspired by Terragni et al. [41]. We choose zero-shot prompting mainly for its simplicity, however, we acknowledge that other approaches, such as few-shot prompting and specific prompts per action, may be considered for future work.

We base our implementation on the UserSimCRS toolkit [2]. The original implementation of *ABUS* is adapted to support LLM-based natural language understanding and generation. *LLM-US* is added as a new user simulator, along with a wrapper to interact with the CRSs. Note that both user simulators use the LLM Mistral NeMo 12b<sup>4</sup> hosted on an Ollama server.<sup>5</sup> For the experiments, we keep the default model options except for the temperature which is set to 0 and deactivate the stream option. This model is chosen for its

<sup>4</sup>Information related to the model is available at: <https://ollama.com/library/mistral-nemo:12b/blobs/b559938ab7a0>

<sup>5</sup><https://ollama.com>

open-source availability and demonstrated performances in natural language generation tasks. Moreover, we release the different prompts used by both simulators.<sup>6</sup>

## 6.2 Synthetic Conversation Generation

For the simulation-based evaluation of the CRSs, we generate 200 synthetic dialogues for each (user simulator, CRS) pair, except for *ABUS-ChatCRS\_OpenDialKG* where the conversation generation failed due to recursion errors. To generate a conversation, the user simulator is initialized with a goal prior to the start of the conversation and then interacts with the CRS with the objective of completing that goal. The goal comprises a set of constraints regarding the movies it is interested in, such as the genre and the year of release, and a set of requests corresponding to information it wants to get on the recommended movies, like its plot and director. An example of a goal is:

$$\text{constraints} = \begin{bmatrix} \text{year} = 2016 \\ \text{actor} = \text{Nicole Kidman} \end{bmatrix} \quad \text{requests} = \begin{bmatrix} \text{plot} \\ \text{director} \\ \text{rating} \end{bmatrix}$$

One advantage of this approach is that it allows for the generation of diverse conversations, as there are many possible combinations of constraints and requests. Furthermore, the user simulator is not looking for any specific target movie, which might get leaked during the conversation. Instead, it considers that there are multiple valid movies satisfying the constraints that may be recommended. This contrasts with the approach taken by [46], where a list of target items is directly provided to the user simulator, a scenario we argue is often unrealistic. Indeed, in typical real-world interactions with a CRS, a user does not have a specific movie in mind but is rather seeking recommendations that satisfy their current preferences.

Following the same approach as in Section 5, we automatically annotate utterances with dialogue acts.

## 6.3 Results

We evaluate CRSs in terms of Reward-per-Dialogue-Length (RDL) and report the results in Table 3. Following the framework proposed by Bernard and Balog [8], we examine the reliability of the user simulators by comparing the results obtained with simulated users (Table 3) against those obtained with real users (Table 2) serving as the reference. The comparison focuses on two aspects: (1) the ability to reproduce the relative ranking of CRSs and (2) the approximation of absolute performance measurements.

For the first aspect, we compute the Kendall’s  $\tau$  correlation between the rankings of CRSs obtained with the simulators and the one based on self-reported user satisfaction. We find correlations of 0.366 for *LLM-US* and 0.143 for *ABUS*, indicating an advantage for *LLM-US*. This is specifically observed for pairwise comparisons of CRSs with the same underlying architecture, where *LLM-US* ranks them in accordance with user satisfaction, except for *UniCRS\_OpenDialKG* and *UniCRS\_ReDial*. However, while the correlations are positive, they are weak, suggesting that the simulators are not yet robust enough to assess a diverse set of CRSs.

**Table 3: Evaluation of CRSs in terms of Reward-per-Dialogue-Length with simulated users. Absolute score differences between measurements obtained with real users and simulated users is shown in parentheses.**

CRS	ABUS	LLM-US
ChatCRS_OpenDialKG	–	0.074 (0.037)
ChatCRS_ReDial	0.121 (0.096)	0.049 (0.024)
BARCOR_ReDial	0.126 (0.102)	0.009 (0.015)
BARCOR_OpenDialKG	0.117 (0.099)	0.011 (0.007)
UniCRS_ReDial	0.079 (0.060)	0.026 (0.007)
KBRD_ReDial	0.084 (0.076)	0.016 (0.008)
CRB-CRS_ReDial	0.070 (0.047)	0.011 (0.012)
UniCRS_OpenDialKG	0.091 (0.085)	0.028 (0.022)
KBRD_OpenDialKG	0.298 (0.294)	0.000 (0.004)
Average abs. diff.	0.107	0.015

Regarding the second aspect, we report the absolute differences between RDL obtained from the conversations with real users and with simulated users in Table 3. The results show that *LLM-US* is better than *ABUS* at approximating the absolute performances obtained with real users, exhibiting average differences of 0.015 and 0.107, respectively.

Based on these observations, we conclude that *LLM-US* outperforms *ABUS* in both reproducing the relative ranking and approximating the absolute performance measurements. While *LLM-US* generally aligns with user satisfaction in ranking CRSs sharing an underlying architecture, it struggles with a more diverse set of systems. This suggests that user simulation, particularly *LLM-US*, presents a viable approach for automating the evaluation of different versions of the same CRS but requires further development to reliably assess broader range of CRSs.

## 7 Discussion and Conclusion

In this work, we addressed the limitations of traditional CRS evaluation by proposing a novel reward-cost framework for measuring user utility and demonstrating its alignment with user satisfaction through two new metrics. Furthermore, we explored the potential of user simulation for automated CRS evaluation, comparing the reliability of different simulation approaches. We now conclude by discussing several crucial aspects for future work in the automatic evaluation of CRSs. These considerations are informed by our examination of current automatic evaluation practices and simulation-based evaluation of CRSs as a solution to alleviate the dependence on real users and human annotators.

**Choice of evaluation metrics.** Different metrics provide different insights and can lead to different conclusions on performance, as our experiments illustrate. We acknowledge that the evaluation of CRSs is still an open and complex challenge, hence, an agreed-upon set of such metrics has not been established yet. We advocate for researchers to prioritize metrics that effectively proxy *utility* for the user as the ultimate goal. To foster the development in this direction, we have presented a general Reward/Cost framework and shown that existing evaluation metrics (Recall@N and success rate) can be understood as specific instantiations within

<sup>6</sup><https://github.com/iai-group/recsys25-crseval>

this framework. We encourage future work to explore novel metrics through the lens of user-perceived reward and interaction cost. While we proposed two initial metrics based on single reward and cost factors, showing positive correlation with user satisfaction, future metrics should incorporate a broader range of factors with dynamically adjusted weights to better capture the nuances of user interactions. This need is underscored by direct user feedback from CRSArena-Dial [9], where some users reported frustration due to, for example, misunderstandings and response time (factors that are not accounted for in our proposed metrics), although the CRS was able to provide satisfactory recommendations.

**Choice of baseline CRSs.** Our results have indicated a large potential headroom based on self-reported user satisfaction. When benchmarking new CRSs, we recommend that researchers consider at least the following two baselines as representatives of two main classes of approaches: (1) a retrieval-based CRS and (2) a LLM-based CRS. The former is conceptually simpler, requires less resources to train and deploy, and is shown to outperform some of the neural CRSs regarding the proposed Reward-per-Dialogue-Length measure. The latter corresponds to the best type of CRSs with regard to user satisfaction according to our results.

**Reliability of user simulators.** Our experiments have shown that the LLM-based simulator, *LLM-US*, is more reliable than the agenda-based one, *ABUS*, at evaluating a diverse set of CRSs. However, *LLM-US* still has shortcomings, especially when it comes to accurately reproducing the relative ranking of a diverse set of CRSs. Therefore, we emphasize that results obtained through simulation-based evaluation necessitate validation by human annotators, as our findings highlight that conclusions can vary significantly depending on the specific user simulator employed.

**Reproducibility.** In practice, a major hurdle to performing simulation-based evaluation is the limited availability and suitability of resources in terms of CRSs and user simulators. To overcome this, it is important to make system implementations publicly available, including a solution for direct interaction (e.g., an API or user interface) and clear setup instructions. Indeed, when performing this study, we have encountered difficulties implementing direct interactions with CRSs due to unclear descriptions of the data format, required annotations, and parameters expected by the system. Although some of the CRS authors provided some support, the technical expertise required to operate the system remains high.

## Limitations

We acknowledge that this work has some limitations. First, our selection of CRSs and user simulators is not exhaustive and is biased towards publicly available, English-language systems used in previous studies. Our analysis in Section 5.3, which examines the correlation between the different utility metrics and self-reported user satisfaction, includes all CRSs, even those with very low user satisfaction, which could potentially skew our observations. However, an additional analysis on a subset of the best-performing systems in terms of satisfaction (*ChatCRS\_\** and *BARCOR\_\**) shows that Kendall’s  $\tau$  correlations with utility metrics remain largely

**Table 4: Kendall’s  $\tau$  correlation between utility metrics and self-reported user satisfaction considering a subset of CRSs (*ChatCRS\_\** and *BARCOR\_\**).**

Metric	Kendall’s $\tau$
Recall@1	0.0
Success rate	0.67
Successful rec. round ratio	-0.34
Reward-per-Dialogue-Length	1.0

consistent (Table 4), with the exception of *SRRR* (which appears to be more sensitive than other metrics). We therefore retained the full set of CRSs to provide a broader perspective on the current state-of-the-art.

Second, our reliance on an LLM for dialogue act annotation may introduce inaccuracies in identifying the acceptance, rejection, and recommendation acts used to compute our proposed utility metrics. We note that the joint annotation of intent and slot values pairs as a multi-label and multi-class classification task is complex and understudied problem [48] (proposed models commonly consider a single dialogue act per utterance). For Mistral NeMo, we observed an area under the ROC curve of 0.673 and 0.683 for user and system intents, respectively, across 77 manually annotated dialogues. While precision scores for Accept (0.61), Reject (0.77), and Recommend (0.96) indicate room for improvement, this approach represents a cost-effective solution for our analysis.

Additionally, some CRSs and user simulators rely on prompts and LLMs that may not be optimal, as different configurations may lead to different results. To mitigate this, the prompts used in this work are inspired and adapted from previous studies [20, 41].

From an ethical perspective, the use of LLMs is associated with well-documented concerns, including the potential to generate harmful content, reinforce biases, and violate privacy, in addition to their environmental impact [7]. We acknowledge that we did not implement explicit moderation or safeguards to address these important issues.

Finally, the evaluation is conducted solely in the movie domain, focusing on specific aspects of the CRSs, which may limit the generalizability of the findings. Despite this, we believe the proposed methodology is broadly applicable for evaluating CRSs in other domains and languages.

## Acknowledgments

This research was supported by the Norwegian Research Center for AI Innovation, NorwAI (Research Council of Norway, project number 309834).

## References

- [1] Zahra Abbasiantaeb, Yifei Yuan, Evangelos Kanoulas, and Mohammad Aliannejadi. 2024. Let the LLMs Talk: Simulating Human-to-Human Conversational QA via Zero-Shot LLM-to-LLM Interactions. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*. 8–17.
- [2] Jafar Afzali, Aleksander Mark Drzewiecki, Krisztian Balog, and Shuo Zhang. 2023. UserSimCRS: A User Simulation Toolkit for Evaluating Conversational Recommender Systems. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining (WSDM '23)*. 1160–1163.
- [3] Leif Azzopardi and Guido Zuccon. 2018. *Economic Models of Interaction*. In *Computational Interaction*. Oxford University Press.

- [4] Leif Azzopardi and Guido Zuccon. 2019. Building Economic Models of Human Computer Interaction. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (CHI EA '19)*. 1–4.
- [5] Krisztian Balog and ChengXiang Zhai. 2024. User Simulation for Evaluating Information Access Systems. *Foundations and Trends in Information Retrieval* 18, 1-2 (2024), 1–261.
- [6] Alejandro Bellogin, Pablo Castells, and Ivan Cantador. 2011. Precision-oriented evaluation of recommender systems: an algorithmic comparison. In *Proceedings of the Fifth ACM Conference on Recommender Systems (RecSys '11)*. 333–336.
- [7] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. 610–623.
- [8] Nolwenn Bernard and Krisztian Balog. 2024. Towards a Formal Characterization of User Simulation Objectives in Conversational Information Access. In *The 10th ACM SIGIR / The 14th International Conference on the Theory of Information Retrieval (ICTIR '24)*.
- [9] Nolwenn Bernard, Hideaki Joko, Faegheh Hasibi, and Krisztian Balog. 2025. CRS Arena: Crowdsourced Benchmarking of Conversational Recommender Systems. In *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining (WSDM '25')*.
- [10] Wanling Cai and Li Chen. 2020. Predicting User Intents and Satisfaction with Dialogue-based Conversational Recommendations. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20)*. 33–42.
- [11] Qibin Chen, Junyang Lin, Yichang Zhang, Ming Ding, Yukuo Cen, Hongxia Yang, and Jie Tang. 2019. Towards Knowledge-Based Recommender Dialog System. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP '19)*. 1803–1813.
- [12] Cheng-Han Chiang and Hung-yi Lee. 2023. Can Large Language Models Be an Alternative to Human Evaluations?. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) (ACL '23)*. 15607–15631.
- [13] Paolo Cremonesi, Yehuda Koren, and Roberto Turrin. 2010. Performance of recommender algorithms on top-n recommendation tasks. In *Proceedings of the Fourth ACM Conference on Recommender Systems (RecSys '10)*. 39–46.
- [14] Jeffrey Dalton, Victor Ajayi, and Richard Main. 2018. Vote Goat: Conversational Movie Recommendation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '18)*. 1285–1288.
- [15] Luke Friedman, Sameer Ahuja, David Allen, Zhenning Tan, Hakim Sidahmed, Changbo Long, Jun Xie, Gabriel Schubiner, Ajay Patel, Harsh Lara, Brian Chu, Zexi Chen, and Manoj Tiwari. 2023. Leveraging Large Language Models in Conversational Recommender Systems. arXiv:2305.07961 [cs.IR]
- [16] Chongming Gao, Wenqiang Lei, Xiangnan He, Maarten de Rijke, and Tat-Seng Chua. 2021. Advances and challenges in conversational recommender systems: A survey. *AI Open* 2 (2021), 100–126.
- [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 & Knowledge Management (CIKM '20)*. 3405–3408.
- [18] Helia Hashemi, Jason Eisner, Corby Rosset, Benjamin Van Durme, and Chris Kedzie. 2024. LLM-Rubric: A Multidimensional, Calibrated Approach to Automated Evaluation of Natural Language Texts. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) (ACL '24)*. 13806–13834.
- [19] Zhankui He, Zhouhang Xie, Rahul Jha, Harald Steck, Dawen Liang, Yesu Feng, Bodhisattwa Prasad Majumder, Nathan Kallus, and Julian McAuley. 2023. Large Language Models as Zero-Shot Conversational Recommenders. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*. 720–730.
- [20] Chen Huang, Peixin Qin, Yang Deng, Wenqiang Lei, Jiancheng Lv, and Tat-Seng Chua. 2024. Concept – An Evaluation Protocol on Conversation Recommender Systems with System-centric and User-centric Factors. arXiv:2404.03304 [cs.CL]
- [21] Dietmar Jannach. 2023. Evaluating conversational recommender systems. *Artificial Intelligence Review* 56, 3 (2023), 2365–2400.
- [22] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. 2021. A Survey on Conversational Recommender Systems. *Comput. Surveys* 54, 5 (2021), 1–36.
- [23] Dan Jurafsky and James H. Martin. 2023. Chatbots and Dialogue Systems. In *Speech and Language Processing* (3rd edition ed.). Chapter 15.
- [24] Wenqiang Lei, Xiangnan He, Yisong Miao, Qingyun Wu, Richang Hong, Min-Yen Kan, and Tat-Seng Chua. 2020. Estimation-Action-Reflection: Towards Deep Interaction Between Conversational and Recommender Systems. In *Proceedings of the 13th International Conference on Web Search and Data Mining (WSDM '20)*. 304–312.
- [25] 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.
- [26] Raymond Li, Samira Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018. Towards Deep Conversational Recommendations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS '18)*. 9748–9758.
- [27] Hsien-chin Lin, Christian Geishauer, Shutong Feng, Nurul Lubis, Carel van Niekerk, Michael Heck, and Milica Gasic. 2022. GenTUS: Simulating User Behaviour and Language in Task-oriented Dialogues with Generative Transformers. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL '22)*. 270–282.
- [28] Hsien-chin Lin, Nurul Lubis, Songbo Hu, Carel van Niekerk, Christian Geishauer, Michael Heck, Shutong Feng, and Milica Gasic. 2021. Domain-independent User Simulation with Transformers for Task-oriented Dialogue Systems. In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL '21)*. 445–456.
- [29] Ahtsham Manzoor and Dietmar Jannach. 2021. Generation-based vs. Retrieval-based Conversational Recommendation: A User-Centric Comparison. In *Proceedings of the 15th ACM Conference on Recommender Systems (RecSys '21)*. 515–520.
- [30] Ahtsham Manzoor and Dietmar Jannach. 2022. Towards retrieval-based conversational recommendation. *Information Systems* 109 (2022), 102083.
- [31] Ahtsham Manzoor, Samuel C. Ziegler, Klaus Maria. Pirker Garcia, and Dietmar Jannach. 2024. ChatGPT as a Conversational Recommender System: A User-Centric Analysis. In *Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization (UMAP '24)*. 267–272.
- [32] Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. OpenDialog: Explainable Conversational Reasoning with Attention-based Walks over Knowledge Graphs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL '19)*. 845–854.
- [33] Arvind Neelakantan, Tao Xu, Raul Puri, Alec Radford, Jesse Michael Han, Jerry Tworek, Qiming Yuan, Nikolas Tezak, Jong Wook Kim, Chris Hallacy, Johannes Heidecke, Pranan Shyam, Boris Power, Tyna Eloundou Nekoul, Girish Sastry, Gretchen Krueger, David Schnurr, Felipe Petroski Such, Kenny Hsu, Madeleine Thompson, Tabarak Khan, Toki Sherbakov, Joanne Jang, Peter Welinder, and Lilian Weng. 2022. Text and Code Embeddings by Contrastive Pre-Training. arXiv:2201.10005 [cs.CL]
- [34] Dhanya Pramod and Prafulla Bafna. 2022. Conversational recommender systems techniques, tools, acceptance, and adoption: A state of the art review. *Expert Systems with Applications* 203 (2022), 117539.
- [35] Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye, and Steve Young. 2007. Agenda-Based User Simulation for Bootstrapping a POMDP Dialogue System. In *Proc. of NAACL '07*. 149–152.
- [36] Ivan Sekulic, Silvia Terragni, Victor Guimaraes, Nghia Khau, Bruna Guedes, Modestas Filipavicius, Andre Ferreira Manso, and Roland Mathis. 2024. Reliable LLM-based User Simulator for Task-Oriented Dialogue Systems. In *Proceedings of the 1st Workshop on Simulating Conversational Intelligence in Chat (SCI-CHAT '24)*. 19–35.
- [37] Harald Steck. 2013. Evaluation of recommendations: rating-prediction and ranking. In *Proceedings of the 7th ACM Conference on Recommender Systems (RecSys '13)*. 213–220.
- [38] Ruixuan Sun, Xinyi Li, Avinash Akella, and Joseph A. Konstan. 2024. Large Language Models as Conversational Movie Recommenders: A User Study. arXiv:2404.19093 [cs.IR]
- [39] 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.
- [40] Yueming Sun and Yi Zhang. 2018. Conversational Recommender System. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '18)*. 235–244.
- [41] Silvia Terragni, Modestas Filipavicius, Nghia Khau, Bruna Guedes, André Manso, and Roland Mathis. 2023. In-Context Learning User Simulators for Task-Oriented Dialog Systems. arXiv:2306.00774 [cs.CL]
- [42] Cynthia A Thompson, Mehmet H Goker, and Pat Langley. 2004. A Personalized System for Conversational Recommendations. *Journal of Artificial Intelligence Research* 21 (2004), 393–428.
- [43] Daniel Valcarce, Alejandro Bellogin, Javier Parapar, and Pablo Castells. 2018. On the robustness and discriminative power of information retrieval metrics for top-n recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18)*. 260–268.
- [44] Maria Vlachou and Craig Macdonald. 2024. What Else Would I Like? A User Simulator using Alternatives for Improved Evaluation of Fashion Conversational Recommendation Systems. arXiv:2401.05783 [cs.IR]
- [45] Ting-Chun Wang, Shang-Yu Su, and Yun-Nung Chen. 2022. BARCOR: Towards A Unified Framework for Conversational Recommendation Systems. arXiv:2203.14257 [cs.CL]

- [46] Xiaolei Wang, Xinyu Tang, Xin Zhao, Jingyuan Wang, and Ji-Rong Wen. 2023. Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP '23)*. 10052–10065.
- [47] Xiaolei Wang, Kun Zhou, Ji-Rong Wen, and Wayne Xin Zhao. 2022. Towards Unified Conversational Recommender Systems via Knowledge-Enhanced Prompt Learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22)*. 1929–1937.
- [48] Henry Weld, Xiaoqi Huang, Siqu Long, Josiah Poon, and Soyeon Caren Han. 2022. A Survey of Joint Intent Detection and Slot Filling Models in Natural Language Understanding. *Comput. Surveys* 55, 8 (2022), 1–38.
- [49] Kerui Xu, Jingxuan Yang, Jun Xu, Sheng Gao, Jun Guo, and Ji-Rong Wen. 2021. Adapting User Preference to Online Feedback in Multi-round Conversational Recommendation. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM '21)*. 364–372.
- [50] Se-eun Yoon, Zhankui He, Jessica Echterhoff, and Julian McAuley. 2024. Evaluating Large Language Models as Generative User Simulators for Conversational Recommendation. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers) (NAACL '24)*. 1490–1504.
- [51] Dianer Yu, Qian Li, Xiangmeng Wang, Qing Li, and Guandong Xu. 2024. Counterfactual Explainable Conversational Recommendation. *IEEE Transactions on Knowledge and Data Engineering* 36, 6 (2024), 2388–2400.
- [52] 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 & Data Mining (KDD '20)*. 1512–1520.
- [53] Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. DIALOGPT : Large-Scale Generative Pre-training for Conversational Response Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations (ACL '20)*. 270–278.
- [54] Kun Zhou, Xiaolei Wang, Yuanhang Zhou, Chenzhan Shang, Yuan Cheng, Wayne Xin Zhao, Yaliang Li, and Ji-Rong Wen. 2021. CRSLab: An Open-Source Toolkit for Building Conversational Recommender System. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations (ACL-IJCNLP '21)*. 185–193.
- [55] Kun Zhou, Wayne Xin Zhao, Shuqing Bian, Yuanhang Zhou, Ji-Rong Wen, and Jingsong Yu. 2020. Improving Conversational Recommender Systems via Knowledge Graph based Semantic Fusion. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '20)*. 1006–1014.
- [56] Yuanhang Zhou, Kun Zhou, Wayne Xin Zhao, Cheng Wang, Peng Jiang, and He Hu. 2022. C<sup>2</sup>-CRS: Coarse-to-Fine Contrastive Learning for Conversational Recommender System. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining (WSDM '22)*. 1488–1496.