

UserSimCRS v2: Simulation-Based Evaluation for Conversational Recommender Systems

Nolwenn Bernard¹ and Krisztian Balog²

¹ TH Köln, Köln, Germany

`nolwenn.bernard@th-koeln.de`

² University of Stavanger, Stavanger, Norway

`krisztian.balog@uis.no`

Abstract. Resources for simulation-based evaluation of conversational recommender systems (CRSs) are scarce. The UserSimCRS toolkit was introduced to address this gap. In this work, we present UserSimCRS v2, a significant upgrade aligning the toolkit with state-of-the-art research. Key extensions include an enhanced agenda-based user simulator, introduction of large language model-based simulators, integration for a wider range of CRSs and datasets, and new evaluation utilities. We demonstrate these extensions in a case study.

Keywords: Conversational Recommender Systems · User simulation · Evaluation

1 Introduction

There is a growing interest in conversational recommender systems (CRSs), which are designed to support users in finding items corresponding to their needs and preferences through a multi-turn dialogue [13]. Despite the progress in the field, CRS evaluation remains a challenging task [4, 5, 13, 17]. Indeed, traditional evaluation methods using offline test collections do not capture the interactive nature of the systems. To address this issue, evaluation methods relying on humans are often employed, yet they suffer from significant limitations, such as their reproducibility, scalability, and cost [18]. Therefore, user simulation has recently been gaining interest to mitigate these limitations [4, 5, 7, 9, 16, 26, 32, 35, 38, 41, 44]. While user simulation can be imperfect, it serves as a valuable preliminary step before human evaluation to refine and make a pre-selection of the best-performing systems, thereby optimizing the use of human resources [1].

Despite this need, there are few open-source resources specifically designed to perform simulation-based evaluation of CRSs. While comprehensive toolkits exist for building CRSs, such as CRSLab [43] and RecWizard [39], they generally lack robust and reusable resources for simulation. This creates a critical gap for researchers, who are often forced to build bespoke evaluation setups from scratch. We also observe that toolkits designed to study task-oriented dialogue systems—a generalization of CRSs to support other tasks than recommendation—include user simulators to perform training and evaluation [21, 31, 45, 46]. However,

these are ill-suited for the recommendation domain, as they are not designed to model crucial recommendation-specific constructs, such as the assessment of recommended items based on preferences from historical interactions. Furthermore, they mostly lack support for standard recommendation benchmark datasets like ReDial [22] and INSPIRED [15].

Recently, Afzali et al. [1] presented UserSimCRS, a toolkit specifically tailored for simulation-based evaluation of CRSs. This original version (v1) provides the implementation of an agenda-based user simulator and basic evaluation metrics. While agenda-based simulators are foundational to the field, the research focus has decisively shifted towards user simulators based on large language models (LLMs), which represent the current state-of-the-art [20, 35–37]. This trend is driven by their advanced natural language understanding capabilities and ability to generate more fluent, diverse, and human-like interactions, moving beyond the often rigid and predictable nature of traditional rule-based approaches, thereby promising substantially greater fidelity in modeling complex user behaviors and preferences.

However, even with a capable simulator, a significant practical challenge remains: interoperability. A simulation-based evaluation setup requires the simulator to communicate with the CRS. This integration is a major barrier to adoption, as CRSs are often implemented in different frameworks and trained on disparate datasets. Facilitating the communication between the simulator and the CRS is essential for enabling reproducible, large-scale comparative studies.

Therefore, the objective of this work is to extend UserSimCRS to provide a more holistic and modernized simulation-based CRS evaluation framework that aligns with the current state-of-the-art in research. These extensions include:

- *Enhanced agenda-based simulator*: We upgrade the classical agenda-based simulator with LLM-based components for dialogue act extraction and natural language generation, and revise the dialogue policy.
- *LLM-based simulators*: We introduce two end-to-end LLM-based user simulators (single-prompt and dual-prompt).
- *Integration with existing CRSs*: We introduce a new communication interface to integrate commonly used CRS models available in CRS Arena [6].
- *Unified data framework*: We provide a unified data format with conversion and LLM-powered augmentation tools for widely-used benchmarks (ReDial [22], INSPIRED [15], and IARD [8]).
- *Advanced conversational quality evaluation*: We expand the toolkit’s limited evaluation metrics by adding a new “LLM-as-a-Judge” utility along with the implementation of recently proposed user-centric utility metrics [5].

We showcase these extensions through a case study in the movie recommendation domain. Specifically, we perform simulation-based evaluation of a selection of CRSs using the different types of user simulators and datasets now supported in UserSimCRS v2. The most recent version of UserSimCRS is available at <https://github.com/iai-group/UserSimCRS>.

2 Related Work

User simulation for the evaluation of interactive information access systems has seen a renewed interest in recent years [4, 7, 26]. Our work is situated within this trend, focusing specifically on conversational recommender systems.

Developing a CRS is a technically challenging task. Therefore, various toolkits have been proposed to lower the technical barrier. For example, CRSLab [43] toolkit supports development with deep neural networks, while FORCE [25] and RecWizard [39] support rule-based and LLM-based CRSs, respectively. These toolkits come with a user interface allowing direct interaction with the CRS and provide limited resources for evaluation. Moreover, offline evaluation commonly focuses on specific components of the CRS, e.g., evaluation of the recommender with mean reciprocal rank or natural language generator with BLEU [43]. This contrasts with platforms dedicated to large-scale *human* evaluation, such as CRS Arena [6], which benchmarks CRSs by having users judge pairwise battles between anonymous systems.

User simulators are commonly built using either a modular or end-to-end architecture. In a modular architecture, different components with specific roles interact with each other to process the incoming utterance and generate a response. For example, the established agenda-based simulator [27] typically comprises three components, as illustrated in Fig. 1a, to transform the incoming utterance to a structured representation (i.e., dialogue acts), decide on the next action to take, and then transform it to a natural response. While in an end-to-end architecture, the user simulator processes the incoming utterance and directly generates the response in natural text (Fig. 1b), e.g., [19, 35]. This architecture is commonly used for user simulators using a (deep) neural network or a large language model.

Simulation-based evaluation of CRSs is still in its early stages, in part due to the complexity of the task and the limited availability of resources [4]. However, some initiatives have recently been proposed to reduce this gap and make simulation-based evaluation more accessible. These include the evaluation methodologies iEvaLM [35] and CONCEPT [16] (built on top of iEvaLM), which share similar objectives to UserSimCRS. They facilitate simulation-based evaluation of CRSs using user simulators based on LLMs. iEvaLM considers objective (e.g., recall) and subjective (e.g., persuasiveness) metrics, while CONCEPT uses system- and user-centric factors, such as reliability, cooperation, and social awareness, to assess a CRS’s performance. However, we note that their implementation exclusively supports CRSs implemented in CRSLab [43] and trained on either ReDial [22] or OpenDialKG [24].

The use of LLMs as surrogates for humans in evaluation, often termed “LLM-as-a-Judge” [42], has become a prevalent trend. This approach is now widely applied across diverse tasks, including relevance assessment in information retrieval [11, 12, 30] and multi-faceted evaluation of CRSs based on detailed scoring rubrics [16, 35]. Despite its widespread adoption, significant open questions remain, including susceptibility to biases and manipulation [2, 3, 33].

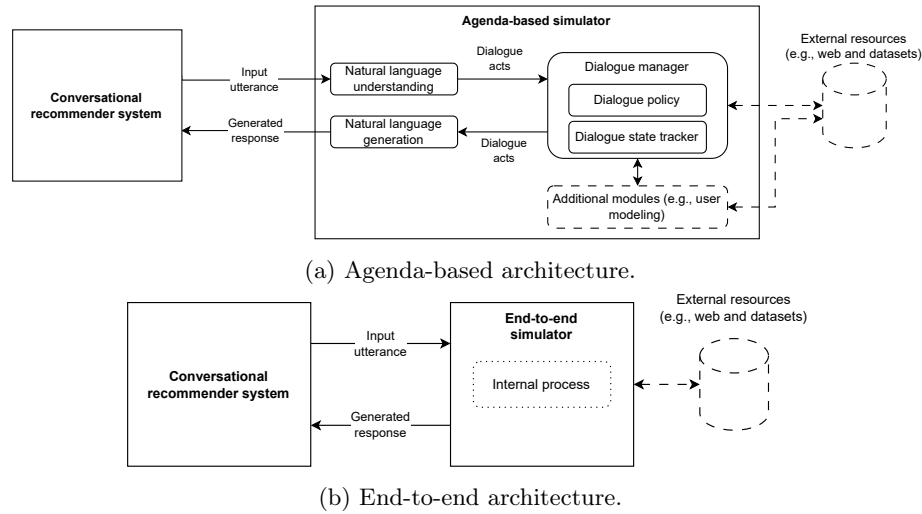


Fig. 1: Agenda-based (Top) and end-to-end (Bottom) architectures for user simulators. The dashed lines represent optional components and data flows.

3 UserSimCRS v1

UserSimCRS [1] is a toolkit facilitating simulation-based evaluation of conversational recommender systems. It is built on top of DialogueKit,³ a library providing dialogue management and evaluation functionalities, and comprises three main modules:

- *Agenda-based user simulator*: The implementation of a modular agenda-based user simulator, which contains the logic to generate user responses based on the dialogue history and the user’s model.
- *User modeling*: Allows the definition of user models covering their preferences, the context, and their persona.
- *Item collection*: Provides a unified representation of the possible items to recommend.

The evaluation process consists of four main steps. First, the experimental setup is defined, including the interaction model, the user population(s), and the evaluation metrics. In the subsequent steps, the user simulator is trained and a set of conversations is generated. Finally, the evaluation metrics (e.g., success rate and average number of turns) are computed. This method permits the relative comparison of different CRSs under the same conditions.

While the toolkit is a valuable asset for the community, we believe that the following limitations hinder its widespread adoption.

³ <https://github.com/iai-group/DialogueKit>

- *Limited selection of user simulators*: The toolkit exclusively provides an agenda-based simulator. While foundational, this approach is becoming outdated and lacks the generative flexibility of LLM-based simulators, which are now prevalent and represent the current state-of-the-art research direction.
- *Limited components for agenda-based simulator*: The original implementation of the agenda-based user simulator is modular; however, it relies on a supervised model for natural language understanding and a template-based model for response generation. This limits the diversity of the responses and requires training data to be effective.
- *Lack of explicitly defined information need*: The preferences and values associated with the elicited slots are randomly sampled, which can lead to incoherent conversations. In theory, the user simulator should be initialized with an information need that guides the conversation and the user responses [27].
- *Lack of support for benchmark datasets*: The toolkit provides a dataset based on the IAI MovieBot agent [14] to create a user simulator, but it is not widely used in the field. Hence, the preparation of popular benchmark datasets is left to the toolkit users which represents a significant burden.
- *Limited selection of evaluation measures*: Evaluation functionalities in DialogueKit are limited to the computation of user satisfaction and average number of turns. The computation of other metrics, especially those related to recommendation and conversational quality, is the responsibility of the toolkit user (i.e., the CRS developer).

The objective of this work is to extend UserSimCRS to address these limitations.

4 UserSimCRS v2 Extensions

The underlying objective of UserSimCRS v2 remains the same as its predecessor: to facilitate the comparison of different conversational recommender systems via simulation-based evaluation. Additionally, it can be used to compare and investigate different user simulators. In this section, we present the extensions introduced in UserSimCRS v2, which are illustrated in Fig. 2. Specifically, we describe the integration of existing benchmark datasets (Section 4.1) the improvements to the agenda-based user simulator (Section 4.2), the introduction of LLM-based user simulators (Section 4.3), the integration of existing CRSs (Section 4.5), and the addition of evaluation utilities and metrics to assess conversation quality and user utility (Section 4.6).

4.1 Unified Data Format and Benchmark Datasets

UserSimCRS v1 provides an initial dataset to create a user simulator, built from interactions with IAI MovieBot [14]; however, this dataset is not widely used in the field. In UserSimCRS v2, we aim to reduce the barrier to adoption by providing a unified data format and conversion tools for the widely-used conversational datasets INSPIRED [15] and ReDial [22]. Additionally, we provide an

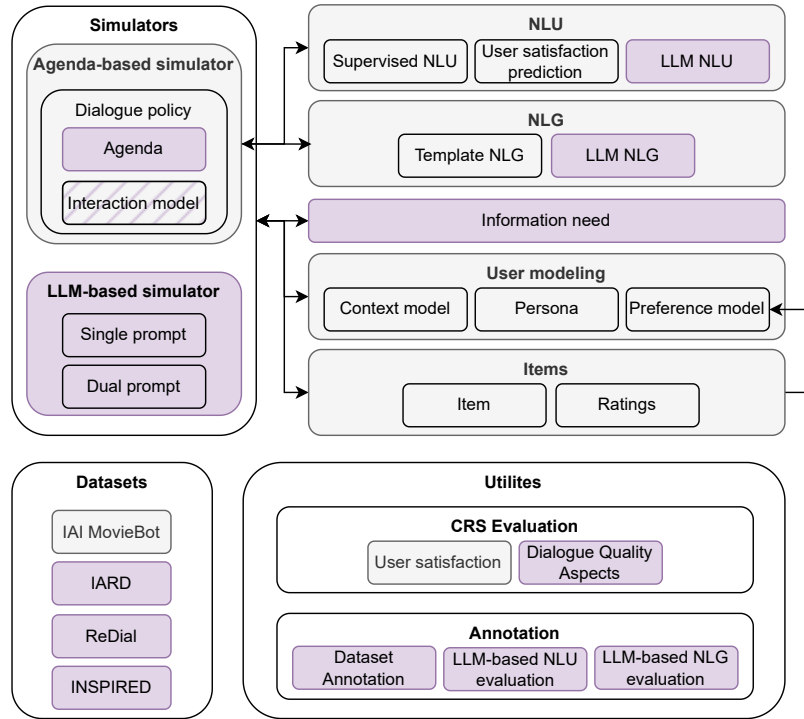


Fig. 2: Overview of UserSimCRS v2 architecture. Grey components are inherited from [1], while hashed and purple components correspond to updated or added components, respectively.

augmented version of IARD [8], a subset of ReDial [22] with intent annotations; see Table 1 for a summary.

The proposed unified data format is based on the format used in DialogueKit, where utterances are annotated with dialogue acts. A dialogue act comprises an intent and optionally a set of associated slot-value pairs. For example, the dialogue acts `Elicit(size)` and `Elicit(color)` are present in the following CRS utterance: “What size and color do you prefer for your new skates?” This allows for a refined representation of utterances and to support more complex interactions than UserSimCRS v1, which only considers a single intent and a list of slot-value pairs per utterance. Furthermore, we introduce the notion of an *information need* that consists of a set of constraints C and requests R as in [27], in addition to target items (i.e., what the user is looking for). For example, the information need of a user looking for the song *Happy* by *Pharrell Williams* and

Table 1: Summary of datasets added in UserSimCRS v2. * indicates augmented version of the original dataset.

Dataset	Domain	# Dialogues	Dialogue annotations
IARD* [8]	Movie	77	Dialogue acts, recommendations accepted
INSPIRED [15]	Movie	1,001	Social strategies, entity mentioned, item feedback
ReDial [22]	Movie	10,006	Movies references, item feedback

who wants to know the release year can be represented as follows:

$$C = \left[\begin{array}{l} \text{genre} = \text{pop} \\ \text{artist} = \text{Pharrell Williams} \end{array} \right]$$

$$R = [\text{release_year} = ?]$$

$$\text{target} = \text{Happy}$$

This information need guides the conversation and the user responses, in addition to informing the user’s preferences and persona. In our example, a response to the CRS utterance eliciting the looked-for genre could be: *“I like pop music, what do you recommend?”*

We note that these benchmark datasets do not necessarily contain all the annotations required to build and train certain user simulators. For example, the utterances in ReDial are not annotated with dialogue acts which are needed to train a user simulator or some components of it (e.g., a supervised natural language understanding model). In such cases, we propose to perform automatic data augmentation. Previous work [28, 40] argues that LLMs are an acceptable alternative to human annotators in data-scarce scenarios. Therefore, in UserSimCRS v2, we provide tools relying on pre-trained LLMs for the tasks of dialogue act annotation and information need extraction. However, when using these tools, the biases and limitations of LLMs should be acknowledged and sanity checks are recommended.

4.2 Enhanced Agenda-based User Simulator

UserSimCRS v1 implements the agenda-based user simulator as a modular system (Fig. 1a) with components for natural language understanding (NLU), dialogue policy,⁴ and natural language generation (NLG). In UserSimCRS v2, we introduce new NLU and NLG components, based on an LLM. They offer ready-to-use alternatives to the original supervised components, which require annotated training data. Furthermore, we revised the dialogue policy component to be guided by the information need and to better reflect the natural flow of recommendation dialogues.

⁴ Referred to as response generation in v1.

Natural Language Understanding. The objective of the NLU component is to extract dialogue acts from incoming CRS utterances. In UserSimCRS v2, we implement a LLM-based NLU component which is initialized with a list of possible intents and slots for the dialogue acts and a prompt that includes a placeholder for the incoming CRS utterance. Dialogue acts extraction is performed by sending the prompt with a given utterance to an LLM and parsing the output based on the following expected format: `intent1(slot="value",slot,...)|intent2()`. We provide a default prompt, built using a few-shot approach with examples taken from the augmented IARD dataset, to extract dialogue acts in the movie domain, along with its corresponding configuration.

Dialogue Policy. The dialogue policy determines the dialogue acts of the next user utterance based on the current agenda and dialogue state. In the agenda-based simulator, this operation is handled by the interaction model. For UserSimCRS v2, the interaction model is revised to align with the new data format that uses dialogue acts and an information need to guide the conversation. In particular, the agenda is now initialized based on the information need, including a disclosure dialogue act for each constraint and an inquiry dialogue act for each request. Furthermore, we modify the agenda update process to better reflect the natural flow of recommendation dialogues, drawing inspiration from Lyu et al. [23]. They observe that an initial elicitation stage typically precedes the first recommendation, after which the user makes inquiries and critiques on the recommended items. Accordingly, our new process explicitly handles four cases: (1) elicitation, (2) recommendation, (3) inquiry, and (4) other. A different type of dialogue act is pushed to the agenda for the first three cases; for example, a disclosure dialogue act is added to the agenda when the CRS elicits information. In the last case, the next dialogue act is either taken from the agenda (if coherent with the current dialogue state) or sampled based on historical data. Note that all the slot-value pairs in the dialogue acts are derived from the information need or from the user’s preferences if not present in the information need.

Natural Language Generation. The NLG component is responsible for generating human-like textual utterance based on the dialogue acts provided by the dialogue policy component. In UserSimCRS v2, we introduce an LLM-based NLG component. It aims to increase the diversity of the generated responses compared to the template-based NLG component from UserSimCRS v1. To generate an utterance, this component uses a prompt that includes placeholders for dialogue acts and, optionally, additional annotations (e.g., emotion). We also provide a default prompt, following a few-shot approach with examples from the augmented IARD dataset, for the movie domain, alongside its corresponding configuration.

4.3 Large Language Model-based User Simulator

The initial version of UserSimCRS only implements agenda-based user simulation. However, other simulation approaches exist, and recently, we observe a growing interest in LLM-based user simulation [37]. To align with this trend and

increase the versatility of the toolkit, UserSimCRS v2 introduces two LLM-based user simulators following different utterance generation strategies.

Single-Prompt Simulator. The single-prompt user simulator produces the next user utterance in an end-to-end manner given a single prompt. The prompt, inspired by Terragni et al. [29], includes the task description, an optional persona, the information need, and the current conversation history. Unlike Terragni et al. [29], our implementation follows a zero-shot approach, i.e., it does not include any examples of conversations. Nevertheless, one could include examples in the task description to follow a few-shot approach.

Dual-Prompt Simulator. The dual-prompt user simulator extends the single-prompt simulator by first using a separate prompt to decide whether to continue the conversation (the “stopping prompt”). If the decision is to continue, it generates the next utterance using the main generation prompt (identical to the one used by the single-prompt simulator); otherwise, it sends a default utterance to stop the conversation. The stopping prompt is designed similarly to the main generation prompt, including the task description (re-framed to focus on taking a binary decision), an optional persona, and the current conversation history.

4.4 Integration with Large Language Models

UserSimCRS v2 introduces several configurable components and user simulators relying on an LLM, all of which assumes the LLM employed is hosted on a server. To manage interactions with the server, all LLM-based modules use a common interface. UserSimCRS v2 provides two initial implementations of this interface: one for OpenAI⁵ and one for an Ollama⁶ server. This design makes the toolkit highly extensible, as additional interfaces for other LLM backends can be easily added through inheritance.

4.5 Integration with Existing CRSs

By design, UserSimCRS considers the CRS as a “black box,” i.e., it does not assume access to its source code or inner workings. The only requirement is to have an interface to interact with the CRS; in practice, this is done using the DialogueKit library. The original version of UserSimCRS provides an interface to communicate with IAI MovieBot [14], a CRS for movie recommendation. In UserSimCRS v2, we introduce a new interface to interact with CRSs available in CRS Arena [6], such as KBRD [10], BARCOR [34], and ChatCRS [35]. This integration reduces the technical burden of using the toolkit, while facilitating experimentation with more diverse and commonly used CRSs.

⁵ <https://openai.com/api/>

⁶ <https://ollama.com/>

4.6 Advanced Evaluation Metrics

Evaluation in UserSimCRS v1 is limited to user satisfaction and average number of turns. However, there are additional measures to assess the performance of CRSs, such as those related to recommendation accuracy and linguistic abilities [17]. In UserSimCRS v2, we provide new evaluation utilities to assess conversation quality using a LLM-based evaluator (“LLM-as-a-Judge”). We also implement novel metrics that measure utility from a user-centric perspective [5].

Conversation Quality Evaluation with LLM-as-a-Judge. Conversation quality is assessed with respect to five aspects inspired by Huang et al. [16]: (1) *Recommendation relevance* measures how closely the recommended items align with the user’s preferences and needs; (2) *Communication style* corresponds to the conciseness and clarity of the responses; (3) *Fluency* is the degree of naturalness of the responses compared to human-generated responses; (4) *Conversational flow* assesses the coherence and consistency of the conversation; and (5) *Overall satisfaction* encapsulates the user’s holistic experience. Each aspect is assigned a score between 1 and 5, with each score defined in a grading rubric.⁷ For example, for *recommendation relevance* a score of 1 represents irrelevant recommendations and a score of 5 represents highly relevant recommendations. We acknowledge that there are open questions regarding the use of LLM-based evaluators, as correlation with human judgments varies across different studies [11].

User-centric Utility Metrics. Bernard and Balog [5] propose metrics to assess the utility of CRSs from a user-centric perspective. We implement Successful Recommendation Round Ratio (SRRR) and Reward-per-Dialogue-Length (RDL), in addition to success rate. SRRR measures the proportion of recommendation rounds resulting in an acceptance, while RDL divides the number of accepted recommendations by dialogue length. The computation of these metrics relies on the underlying intent of user utterances in a conversation, that are annotated either manually or automatically using an NLU model.

5 Use Case: Movie Recommendation

To demonstrate the use of different user simulators and CRSs supported in UserSimCRS v2, we present a case study in the movie recommendation domain.

User Simulators. We instantiate the different types of user simulators supported in UserSimCRS v2—agenda-based (ABUS), single/dual-prompt LLM-based (LLM-SP/DP)—using different datasets. Depending on the type of the simulator and the dataset, we prepare the item collection, define an interaction model, create prompts, and annotate a sample of dialogues with dialogue acts.

⁷ The grading rubrics are available at: https://github.com/iai-group/UserSimCRS/blob/main/scripts/evaluation/rubrics/quality_rubrics.py

Table 2: Evaluation of selected CRSs with different user simulators. Results reported correspond to average scores over 100 synthetic dialogues.

CRS	User satisfaction			Fluency			Rec. relevance		
	ABUS	LLM-SP	LLM-DP	ABUS	LLM-SP	LLM-DP	ABUS	LLM-SP	LLM-DP
<i>Dataset: MovieBot</i>									
BARCOR_OpenDialKG	1.9 ±0.6	1.9 ±0.6	1.9 ±0.6	2.8 ±0.5	2.6 ±0.6	3.0 ±0.5	1.4 ±0.6	1.5 ±0.9	2.0 ±0.9
BARCOR_ReDial	2.0 ±0.5	1.8 ±0.5	2.0 ±0.6	3.3 ±0.8	2.6 ±0.6	2.9 ±0.6	2.4 ±1.6	1.5 ±0.9	1.9 ±0.9
KBRD_ReDial	1.9 ±0.4	1.9 ±0.4	1.9 ±0.4	2.5 ±0.6	2.4 ±0.6	2.6 ±0.6	1.2 ±0.6	1.3 ±0.7	1.5 ±0.8
UniCRS_OpenDialKG	2.3 ±0.8	1.8 ±0.6	1.9 ±0.6	2.4 ±0.5	2.5 ±0.6	2.8 ±0.5	1.0 ±0.2	1.5 ±1.0	1.5 ±0.7
IAI MovieBot	1.5 ±0.7	1.7 ±0.8	2.5 ±0.7	3.4 ±0.6	3.1 ±0.7	3.3 ±1.0	2.1 ±1.1	2.3 ±1.4	2.9 ±1.7
<i>Dataset: INSPIRED</i>									
BARCOR_OpenDialKG	2.2 ±0.7	2.0 ±0.6	2.1 ±0.6	2.6 ±0.6	2.7 ±0.5	3.0 ±0.6	1.6 ±0.6	1.4 ±0.6	1.9 ±1.0
BARCOR_ReDial	2.1 ±0.5	2.0 ±0.5	2.0 ±0.5	3.5 ±0.5	2.6 ±0.6	2.9 ±0.6	2.3 ±0.8	1.4 ±0.9	1.6 ±0.9
KBRD_ReDial	2.0 ±0.3	2.0 ±0.5	2.1 ±0.6	3.3 ±0.6	2.3 ±0.5	2.6 ±0.6	1.9 ±0.9	1.2 ±0.6	1.4 ±0.7
UniCRS_OpenDialKG	1.9 ±0.8	2.1 ±0.6	2.1 ±0.6	2.9 ±0.5	2.4 ±0.5	2.7 ±0.5	1.3 ±0.6	1.1 ±0.4	1.5 ±0.7
IAI MovieBot	1.4 ±0.8	2.2 ±0.8	2.2 ±0.8	3.3 ±0.5	2.9 ±1.0	3.3 ±1.0	2.1 ±0.8	2.7 ±1.9	3.1 ±1.8
<i>Dataset: ReDial</i>									
BARCOR_OpenDialKG	2.2 ±0.8	1.9 ±0.8	2.1 ±0.7	3.0 ±0.6	2.6 ±0.5	3.0 ±0.5	1.9 ±0.8	1.4 ±0.6	1.8 ±0.9
BARCOR_ReDial	2.0 ±0.1	2.2 ±0.6	2.3 ±0.8	3.4 ±0.6	2.4 ±0.5	2.4 ±0.6	2.3 ±0.9	1.2 ±0.6	1.3 ±0.6
KBRD_ReDial	1.9 ±0.6	2.1 ±0.5	2.1 ±0.8	3.0 ±0.5	2.1 ±0.4	2.3 ±0.5	1.7 ±0.8	1.0 ±0.1	1.2 ±0.6
UniCRS_OpenDialKG	2.0 ±0.8	1.9 ±0.8	1.9 ±0.8	2.7 ±0.5	2.3 ±0.5	2.8 ±0.6	1.3 ±0.5	1.1 ±0.3	1.5 ±0.9
IAI MovieBot	2.0 ±0.2	2.3 ±0.9	2.2 ±0.9	4.0 ±0.2	2.9 ±1.0	3.2 ±1.0	4.7 ±0.8	2.4 ±1.7	2.8 ±1.6

Conversational Recommender Systems. We consider four conversational recommender systems from the CRS Arena [6]: BARCOR_OpenDialKG, BARCOR_ReDial, KBRD_ReDial, and UniCRS_OpenDialKG. Additionally, we include IAI MovieBot, as it was already supported in UserSimCRS v1.

Evaluation. We evaluate the selected CRSs using the different user simulators. This process consists of generating 100 synthetic dialogues for each pair of user simulator and CRS, and computing evaluation metrics over these dialogues. In Table 2, we report user satisfaction, fluency, and recommendation relevance.⁸ The results confirm that, in general, CRS performance remains a significant challenge, with most scores below 3 on a 5-point scale. This is inline with observations reported in [5, 6]. However, a closer look at the table reveals some interesting findings.

First, the simulators demonstrate significant disagreement on system ranking. For instance, on the MovieBot dataset, ABUS ranks IAI MovieBot last for user satisfaction (1.5), while LLM-DP ranks it first (2.5). Furthermore, even when simulators agree on the top-performing system (e.g., IAI MovieBot for Recommendation Relevance on ReDial), they often diverge on the magnitude of its performance (a 4.7 from ABUS vs. 2.4/2.8 from the LLM simulators).

The simulators also show distinct characteristics. ABUS appears to be the “most opinionated,” assigning both the highest score in the table (4.7) and some of the lowest (1.0). The LLM simulators, by contrast, operate within a more compressed range, rarely awarding scores above 3.3. Moreover, the LLM simu-

⁸ Configurations and synthetic dialogues are available in the repository.

lators are not interchangeable. Despite producing similar results at times, LLM-DP consistently yields equal or higher average scores than LLM-SP on the INSPIRED dataset across all three metrics.

Finally, CRS performance itself is highly dataset-dependent. Some systems show extreme variability; for example, IAI MovieBot’s recommendation relevance (via ABUS) peaks at 4.7 on ReDial—the highest score in the table—but drops to a mediocre 2.1 on the other datasets. In contrast, other systems demonstrate high consistency. BARCOR_ReDial’s fluency score from ABUS, for instance, remains stable and high across all three datasets (3.3, 3.5, 3.4).

Enabled Research Directions. Considering the results and the synthetic dialogues generated, we can envision different research questions that UserSimCRS v2 can help investigate. For example, these include studying the influence of different user simulator types (e.g., agenda-based vs. LLM-based) and configurations (e.g., different LLMs, different prompts, datasets, and interaction models) on the evaluation outcomes, comparing the performance of different CRSs, and analyzing the synthetic dialogues with regard to their discourse structure, characteristics, or breakdowns. These are merely illustrative of the much broader scope of research that UserSimCRS v2 enables; however, the focus of this paper remains on the presentation of the toolkit itself.

6 Conclusion

In this work, we presented an upgraded version of UserSimCRS, designed to facilitate a more comprehensive and flexible comparison of user simulators and conversational recommender systems through simulation-based evaluation. This new version aligns the toolkit with current research by enhancing the classic agenda-based user simulator with LLM-powered components, introducing LLM-based simulators, adding new automatic evaluation utilities, and providing support for a wider range of CRSs and datasets.

Future work will focus on enhancing the integration of user modeling components to better support the simulation of diverse user populations and exploring novel evaluation metrics. Furthermore, we aim to continue reducing the technical entry barrier to encourage widespread adoption of the toolkit by the community, thereby fostering research on user simulation.

Acknowledgments. This research was partially supported by the Norwegian Research Center for AI Innovation, NorwAI (Research Council of Norway, project number 309834), and by the German Federal Ministry of Education and Research and Joint Science Conference under the P_{LAN}_CV project (reference number 03FHP109).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

Bibliography

- [1] Afzali, J., Drzewiecki, A.M., Balog, K., Zhang, S.: UserSimCRS: A user simulation toolkit for evaluating conversational recommender systems. In: Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining, pp. 1160–1163, WSDM '23 (2023)
- [2] Alaofi, M., Thomas, P., Scholer, F., Sanderson, M.: LLMs can be fooled into labelling a document as relevant. In: Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region, p. 32–41, SIGIR-AP '24 (2024)
- [3] Balog, K., Metzler, D., Qin, Z.: Rankers, judges, and assistants: Towards understanding the interplay of LLMs in information retrieval evaluation. In: Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 3865–3875, SIGIR '25 (2025)
- [4] Balog, K., Zhai, C.: User simulation for evaluating information access systems. *Foundations and Trends in Information Retrieval* **18**(1-2), 1–261 (2024), ISSN 1554-0669
- [5] Bernard, N., Balog, K.: 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, pp. 261–271, SIGIR-AP 2025 (2025)
- [6] Bernard, N., Joko, H., Hasibi, F., Balog, K.: CRS Arena: Crowdsourced benchmarking of conversational recommender systems. In: Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining, pp. 1028–1031, WSDM '25 (2025)
- [7] Breuer, T., Kreutz, C.K., Fuhr, N., Balog, K., Schaer, P., Bernard, N., Frommholz, I., Gohsen, M., Ji, K., Jones, G.J.F., Keller, J., Liu, J., Mladenov, M., Pasi, G., Trippas, J., Wang, X., Zerhoudi, S., Zhai, C.: Report on the workshop on simulations for information access (Sim4IA 2024) at SIGIR 2024. *SIGIR Forum* **58**(2), 1–14 (2025)
- [8] Cai, W., Chen, L.: Predicting user intents and satisfaction with dialogue-based conversational recommendations. In: Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization, pp. 33–42, UMAP '20 (2020)
- [9] Chen, L., Dai, Q., Zhang, Z., Feng, X., Zhang, M., Tang, P., Chen, X., Zhu, Y., Dong, Z.: RecUserSim: A realistic and diverse user simulator for evaluating conversational recommender systems. In: Companion Proceedings of the ACM on Web Conference 2025, pp. 133–142, WWW '25 (2025)
- [10] Chen, Q., Lin, J., Zhang, Y., Ding, M., Cen, Y., Yang, H., Tang, J.: 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, pp. 1803–1813, EMNLP-IJCNLP '19 (2019)
- [11] Dietz, L., Zendel, O., Bailey, P., Clarke, C.L.A., Cotterill, E., Dalton, J., Hasibi, F., Sanderson, M., Craswell, N.: Principles and guidelines for the use of LLM judges. In: Proceedings of the 2025 International ACM SIGIR Conference on Innovative Concepts and Theories in Information Retrieval, pp. 218–229, ICTIR '25 (2025)
- [12] Faggioli, G., Dietz, L., Clarke, C.L.A., Demartini, G., Hagen, M., Hauff, C., Kando, N., Kanoulas, E., Potthast, M., Stein, B., Wachsmuth, H.: Perspectives on large language models for relevance judgment. In: Proceedings of the 2023 ACM

- SIGIR International Conference on Theory of Information Retrieval, p. 39–50, ICTIR '23 (2023)
- [13] Gao, C., Lei, W., He, X., de Rijke, M., Chua, T.S.: Advances and challenges in conversational recommender systems: A survey. *AI Open* **2**, 100–126 (2021)
- [14] Habib, J., Zhang, S., Balog, K.: IAI MovieBot: A conversational movie recommender system. In: Proceedings of the 29th ACM International Conference on Information & Knowledge Management, pp. 3405–3408, CIKM '20 (2020)
- [15] Hayati, S.A., Kang, D., Zhu, Q., Shi, W., Yu, Z.: INSPIRED: Toward sociable recommendation dialog systems. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, pp. 8142–8152, EMNLP '20 (2020)
- [16] Huang, C., Qin, P., Deng, Y., Lei, W., Lv, J., Chua, T.S.: Concept – An evaluation protocol on conversation recommender systems with system-centric and user-centric factors. arXiv **cs.CL/2404.03304** (2024)
- [17] Jannach, D.: Evaluating conversational recommender systems. *Artificial Intelligence Review* **56**(3), 2365–2400 (2023)
- [18] Kelly, D.: Methods for evaluating interactive information retrieval systems with users. *Foundations and Trends® in Information Retrieval* **3**(1-2), 1–224 (2009)
- [19] Kim, M., Kim, M., Kim, H., Kwak, B.w., Kang, S., Yu, Y., Yeo, J., Lee, D.: Pearl: A review-driven persona-knowledge grounded conversational recommendation dataset. In: Findings of the Association for Computational Linguistics: ACL 2024, pp. 1105–1120 (2024)
- [20] Kim, S., Kim, T., Seo, K., Yeo, J., Lee, D.: Stop playing the guessing game! Target-free user simulation for evaluating conversational recommender systems. arXiv **cs.IR/2411.16160** (2024)
- [21] Lee, S., Zhu, Q., Takanobu, R., Zhang, Z., Zhang, Y., Li, X., Li, J., Peng, B., Li, X., Huang, M., Gao, J.: ConvLab: Multi-domain end-to-end dialog system platform. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, pp. 64–69, ACL '19 (2019)
- [22] Li, R., Kahou, S., Schulz, H., Michalski, V., Charlin, L., Pal, C.: Towards deep conversational recommendations. In: Proceedings of the 32nd International Conference on Neural Information Processing Systems, pp. 9748–9758, NIPS '18 (2018)
- [23] Lyu, S., Rana, A., Sanner, S., Bouadjenek, M.R.: A workflow analysis of context-driven conversational recommendation. In: Proceedings of the Web Conference 2021, pp. 866–877, WWW '21 (2021)
- [24] Moon, S., Shah, P., Kumar, A., Subba, R.: OpenDialKG: Explainable conversational reasoning with attention-based walks over knowledge graphs. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 845–854, ACL '19 (2019)
- [25] Quan, J., Wei, Z., Gan, Q., Yao, J., Lu, J., Dong, Y., Liu, Y., Zeng, Y., Zhang, C., Li, Y., Hu, H., He, Y., Yang, Y., Jiang, D.: FORCE: A framework of rule-based conversational recommender system. In: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 13215–13217, AAAI '22 (2022)
- [26] Schaer, P., Kreutz, C.K., Balog, K., Breuer, T., Kruff, A.K.: Second SIGIR workshop on simulations for information access (Sim4IA 2025). In: Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 4172–4175, SIGIR '25 (2025)
- [27] Schatzmann, J., Thomson, B., Weilhammer, K., Ye, H., Young, S.: Agenda-based user simulation for bootstrapping a POMDP dialogue system. In: *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Companion Volume, Short Papers*, pp. 149–152, NAACL '07 (2007)

- [28] Su, H., Ye, J.: Large language models for automating fine-grained speech act annotation: A critical evaluation of GPT-4o and DeepSeek. *Corpus Pragmatics* (2005)
- [29] Terragni, S., Filipavicius, M., Khau, N., Guedes, B., Manso, A., Mathis, R.: In-context learning user simulators for task-oriented dialog systems. *arXiv cs.CL/2306.00774* (2023)
- [30] Thomas, P., Spielman, S., Craswell, N., Mitra, B.: Large language models can accurately predict searcher preferences. In: *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1930–1940, *SIGIR '24* (2024)
- [31] Ultes, S., Rojas-Barahona, L.M., Su, P.H., Vandyke, D., Kim, D., Casanueva, I., Budzianowski, P., Mrkšić, N., Wen, T.H., Gašić, M., Young, S.: PyDial: A multi-domain statistical dialogue system toolkit. In: *Proceedings of ACL 2017, System Demonstrations*, pp. 73–78, *ACL '17* (2017)
- [32] Vlachou, M.: Fashion-AlterEval: A dataset for improved evaluation of conversational recommendation systems with alternative relevant items. In: *Proceedings of the Nineteenth ACM Conference on Recommender Systems*, pp. 755–763, *RecSys '25* (2025)
- [33] Wang, P., Li, L., Chen, L., Cai, Z., Zhu, D., Lin, B., Cao, Y., Kong, L., Liu, Q., Liu, T., Sui, Z.: Large language models are not fair evaluators. In: *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 9440–9450, *ACL '24* (Aug 2024)
- [34] Wang, T.C., Su, S.Y., Chen, Y.N.: BARCOR: Towards a unified framework for conversational recommendation systems. *arXiv cs.CL/2203.14257* (2022)
- [35] Wang, X., Tang, X., Zhao, X., Wang, J., Wen, J.R.: 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*, pp. 10052–10065, *EMNLP '23* (2023)
- [36] Yoon, S.e., He, Z., Echterhoff, J., McAuley, J.: 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)*, pp. 1490–1504, *NAACL '24* (2024)
- [37] Zhang, H., Zhao, X., Chen, J., Guo, J.: A literature review on simulation in conversational recommender systems. *arXiv cs.HC/2506.20291* (2025)
- [38] Zhang, S., Balog, K.: Evaluating conversational recommender systems via user simulation. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1512–1520, *KDD '20* (2020)
- [39] Zhang, Z., Laud, T., He, Z., Chen, X., Liu, X., Xie, Z., McAuley, J., He, Z.: RecWizard: A toolkit for conversational recommendation with modular, portable models and interactive user interface. In: *Proceedings of the 38th Annual AAAI Conference on Artificial Intelligence, AAAI '24* (2024)
- [40] Zhangwenbo, Z., Yuhan, W.: Act2P: LLM-driven online dialogue act classification for power analysis. In: *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 20494–20504, *ACL '25* (2025)
- [41] Zhao, X., Deng, Y., Wang, W., lin, H., Cheng, H., Zhang, R., Ng, S.K., Chua, T.S.: Exploring the impact of personality traits on conversational recommender systems: A simulation with large language models. *arXiv cs.CL/2504.12313* (2025)

- [42] Zheng, L., Chiang, W.L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E.P., Zhang, H., Gonzalez, J.E., Stoica, I.: Judging LLM-as-a-judge with MT-bench and Chatbot Arena. In: Proceedings of the 37th International Conference on Neural Information Processing Systems, NeurIPS '23 (2023)
- [43] Zhou, K., Wang, X., Zhou, Y., Shang, C., Cheng, Y., Zhao, W.X., Li, Y., Wen, J.R.: 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, pp. 185–193, ACL-IJCNLP '21 (2021)
- [44] Zhu, L., Huang, X., Sang, J.: A LLM-based controllable, scalable, human-involved user simulator framework for conversational recommender systems. In: Proceedings of the ACM on Web Conference 2025, pp. 4653–4661, WWW '25 (2025)
- [45] Zhu, Q., Geishauser, C., Lin, H.c., van Niekerk, C., Peng, B., Zhang, Z., Feng, S., Heck, M., Lubis, N., Wan, D., Zhu, X., Gao, J., Gasic, M., Huang, M.: ConvLab-3: A flexible dialogue system toolkit based on a unified data format. In: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pp. 106–123, EMNLP '23 (2023)
- [46] Zhu, Q., Zhang, Z., Fang, Y., Li, X., Takanobu, R., Li, J., Peng, B., Gao, J., Zhu, X., Huang, M.: ConvLab-2: An open-source toolkit for building, evaluating, and diagnosing dialogue systems. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, pp. 142–149, ACL '20 (2020)