



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/245022/>

Version: Accepted Version

Proceedings Paper:

Ramos, J., Xia, F., Wang, X. et al. (2026) Interplay: training independent simulators for reference-free conversational recommendation. In: Campos, R., Jatowt, A., Lan, Y., Aliannejadi, M., Bauer, C., MacAvaney, S., Anand, A., Ren, Z., Verberne, S., Bai, N. and Mansoury, M., (eds.) Advances in Information Retrieval: 48th European Conference on Information Retrieval, ECIR 2026, Delft, The Netherlands, March 29 – April 2, 2026, Proceedings, Part I. 48th European Conference on Information Retrieval, ECIR 2026, 29 Mar - 02 Apr 2026, Delft, The Netherlands. Lecture Notes in Computer Science, LNCS 16483. Springer Nature Switzerland, pp. 629-644. ISBN: 9783032212887. ISSN: 0302-9743. EISSN: 1611-3349.

https://doi.org/10.1007/978-3-032-21289-4_40

© 2026 The Authors. Except as otherwise noted, this author-accepted version of a conference paper published in Advances in Information Retrieval: 48th European Conference on Information Retrieval, ECIR 2026, Delft, The Netherlands, March 29 – April 2, 2026, Proceedings, Part I is made available via the University of Sheffield Research Publications and Copyright Policy under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Interplay: Training Independent Simulators for Reference-Free Conversational Recommendation

Jerome Ramos¹, Feng Xia², Xi Wang², Shubham Chatterjee³, Xiao Fu¹,
Hossein A. Rahmani¹, and Aldo Lipani¹

¹ University College London, London, UK

{jerome.ramos.20,xiao.fu.20,hossein.rahmani.22,aldo.lipani}@ucl.ac.uk

² University of Sheffield, Sheffield, UK

{fxia8,xi.wang}@sheffield.ac.uk

³ Missouri University of Science and Technology, Rolla, USA

shubham.chatterjee@mst.edu

Abstract. Training conversational recommender systems (CRS) requires extensive dialogue data, which is challenging to collect at scale. To address this, researchers have used simulated user-recommender conversations. Traditional simulation approaches often utilize a single large language model (LLM) that generates entire conversations with prior knowledge of the target items, leading to scripted and artificial dialogues. We propose a reference-free simulation framework that trains two independent LLMs, one as the user and one as the conversational recommender. These models interact in real-time without access to predetermined target items, but preference summaries and target attributes, enabling the recommender to genuinely infer user preferences through dialogue. This approach produces more realistic and diverse conversations that closely mirror authentic human-AI interactions. Our reference-free simulators match or exceed existing methods in quality, while offering a scalable solution for generating high-quality conversational recommendation data without constraining conversations to pre-defined target items. We conduct both quantitative and human evaluations to confirm the effectiveness of our reference-free approach.

1 Introduction

Conversational Recommender Systems (CRS) are mixed-initiative dialogue systems that offer personalized, context-sensitive recommendations through natural language conversations [22]. Unlike traditional systems that rely on static user-item interactions, CRS allow dynamic, interactive feedback from users and enable both parties to guide the conversation.

The development of CRS depends on rich conversational data, but collecting human annotated conversations is costly and limited in quality. Previous work highlights key issues in crowd-sourced CRS datasets: (1) **Generic utterances:** Due to lack of genuine preferences, workers often say vague statements like “I like most genres,” limiting personalization [12]; (2) **Lack of context:** Qiu et al. [20]

showed that crowd-sourced conversations often lack depth and context, making it difficult for CRS to accurately model user preferences and provide relevant suggestions; and (3) **Limited domain expertise**: Qiu et al. [20] noted that crowd-workers often lack domain expertise, resulting in weak recommendations and poor explanations. Joko et al. [10] further showed that even well-designed crowd data lacks multi-session realism and is hard to scale without sacrificing authenticity. Human-written conversations are also costly and labor-intensive to collect at scale.

The growing popularity of large language models (LLMs) has led to the use of LLM simulators to generate authentic and context-rich conversations that are critical for training CRS. Zhou et al. [30] found that topic-guided dialogues improve recommendations, underscoring the need for synthetic data that mimic natural interactions. Recent approaches use LLMs like GPT-4o to generate CRS conversations [12, 17, 31], but generally rely on reference-dependent generation, feeding target items to simulators in advance. This causes: (1) data leakage (the simulators know the answer beforehand) and (2) rigid dialogues (exploration is limited, reducing realism). Additionally, using general-purpose LLMs as static role-players for both conversational parties is problematic [13]. These models are designed as helpful assistants rather than realistic user simulators with diverse and sometimes inconsistent preferences. The resulting conversations often lack the natural exploration, uncertainty, and flexibility that characterize genuine human-recommender interactions.




Fig. 1. Unlike prior work that uses templated dialogues and target-aware static LLMs, our method independently trains user and recommender simulators without predetermined items or actions, enabling diverse, realistic interactions.

In this chapter, we address these limitations through an interplay approach that trains two specialized models to interact independently. As depicted in Figure 1, our method fine-tunes separate user and recommender simulators on

existing CRS data using a reference-free framework. Crucially, neither simulator has access to target items during generation. The user simulator operates with preference summaries and attribute descriptions, providing realistic feedback without knowing specific targets. The recommender simulator generates contextually appropriate recommendations based solely on the evolving conversation. This design enables genuine exploration where recommendations emerge naturally from dialogue rather than predetermined scripts.

Our approach offers several advantages over existing methods. The reference-free design eliminates data leakage by ensuring that simulators discover information through conversation rather than having it predetermined for them. Independent training of specialized models creates more realistic user and recommender behaviors than using general-purpose LLMs for both roles. The interactive nature allows flexible conversations that can naturally evolve in multiple directions, better reflecting the complexity of real recommendation scenarios. Through comprehensive evaluations, we demonstrate that our simulators generate conversations comparable in quality to existing approaches while addressing the methodological limitations that have constrained previous work. The resulting synthetic dialogues exhibit natural exploration patterns, genuine uncertainty, and flexibility that characterize human conversations.

We make the following contributions in this chapter:

1. We propose a reference-free simulation framework that removes the oracle problem by training independent user and recommender simulators without access to the target item, allowing for the exploration of realistic preferences and flexible recommendation behavior.
2. We show that simulators trained on target attributes, not items, outperform larger general models: Our 8B recommender achieves $15\times$ higher accuracy than 32B baselines while preserving strong response diversity.
3. Through a qualitative user study, we show that independent simulators match reference-dependent methods in dialogue quality across six metrics, while using smaller open-source models for a more scalable and efficient solution.

2 Related Works

Conversational Recommender Systems. CRS integrate dialogue interfaces with personalized recommendation engines, allowing users to request and receive item suggestions through natural language interaction [7, 22]. Unlike one-shot recommenders, CRS conduct multi-turn conversations to elicit user preferences, clarify requirements, and offer explanations [7]. Recent advances have shifted from modular pipelines [11] to end-to-end models [15] and prompt-based approaches that achieve strong zero-shot performance [4, 18]. However, these systems face challenges in preference modeling, with methods such as CORAL that use LLMs to handle contrasting user preferences [14].

User Simulators. User simulators enable scalable CRS training and evaluation without human participants [28]. Early simulator techniques were rule-based but

too rigid [11], while recent approaches leverage LLMs with persona descriptions for more realistic responses. However, a critical limitation is reference dependency: many methods provide simulators with target items, creating artificial constraints [27]. Frameworks like CSHI balance realism and consistency through human-in-the-loop guidance [32]. Our work addresses these issues through independent simulator training.

Dialogue Data Collection and Generation. CRS research faces a trade-off between data quality and scalability. Human-centric methods produce high quality but are expensive [28], while crowdsourcing offers scale but suffers from generic responses and quality issues [12]. LLM-based generation has emerged as a solution, using techniques such as template-driven approaches [17] and active synthesis [23]. However, maintaining reference-independence while ensuring diversity remains challenging.

Interactive Simulation and Evaluation Challenges. Recent multi-agent approaches in CRS use independent user and recommender simulators to foster natural interactions. Methods like MACRS [3] and Collaborative Self-Play [2] leverage multiple LLMs with distinct roles, while reinforcement learning-based systems such as CRSIRL [1] and AMIE [25] explore self-play and intrinsic reward modeling. However, true agent independence remains unexplored. Simultaneously, evaluation challenges persist: datasets like ReDial [16] and OpenDialKG [19] suffer from data leakage, inflating performance through exposed targets [27]. Newer methods, such as FACE [9], aim to assess the quality of the dialogue without references, addressing the disconnect between traditional metrics and human judgment [28] and highlighting the need for reference-free, target-agnostic evaluation.

3 Methodology

We propose **Interplay** – a framework that trains independent user and recommender simulators to interact without predetermined outcomes, allowing genuine exploration and flexible conversation dynamics that mirror real-world recommendation scenarios.

3.1 The Core Insight

The core challenge in simulating conversational recommendation lies in replicating the natural tension of real interactions: users bring preferences but stay open to discovering unexpected options, while recommenders must explore needs without knowing what will click. This tension is lost when a single model plays both roles with shared knowledge. Our approach addresses this challenge through two key innovations. First, we train completely independent simulators, each optimized solely for its respective role, ensuring that natural behavioral differences emerge rather than being artificially programmed. Second, we remove reference dependence by using target attributes instead of items, prompting simulators to engage in genuine discovery over scripted responses.

3.2 Breaking the Single-Target Constraint

Prior user simulation methods are fundamentally rigid – they condition responses on fixed target movies, making simulated users reject strong recommendations simply because they do not match a preset title [27]. This creates an artificial constraint that real users never face. When someone seeks “a romantic comedy for date night,” they usually remain open to various films that meet those criteria rather than focusing on a single predetermined choice. Our user simulator captures natural flexibility by reasoning over target attributes instead of fixed items. Rather than knowing “that the user wants Inception,” the simulator is trained to understand preferences like “a mind-bending thriller with complex plotting.” This change allows the user simulator to accept a range of valid recommendations, freeing conversations from the constraint of converging on a single outcome. As a result, dialogues unfold more organically, driven by authentic attribute alignment, rather than scripted goals.

3.3 Persona Architecture for Authentic User Behavior

To produce realistic interactions, our user simulator integrates three persona components that capture diverse aspects of user context while preserving flexibility. The *general preferences component* offers a natural language summary of the user’s broad tastes, generated by analyzing a subset of their viewing history. This summary grounds the simulator in a consistent personality across conversations while still allowing variation in specific choices. The *movie history component* includes three previously watched movies with corresponding reviews, randomly sampled from the user’s profile to provide concrete examples of their tastes. These serve as conversational anchors, allowing the recommender to reference and build on specific, discussable preferences. The three random movies were originally sampled with the PEARL dataset and we reuse the same movies in order to maintain the same context and avoid data leakage. The *target attributes component* then introduces preference cues extracted from a target movie’s review, without revealing the movie title. This preserves authentic user intent while removing reference dependence: the simulator knows what qualities to look for, but not which item should satisfy them. This architecture enables our user simulator to maintain a consistent personality while remaining genuinely open to discovery, accurately mirroring how real users approach recommendation conversations with both clear preferences and openness to surprise.

3.4 Recommender Without Oracle Knowledge

Our recommender simulator learns to handle preference elicitation and recommendation generation through fluid conversations, adapting dynamically to evolving user needs rather than following rigid patterns. We cast the recommender task as contextual text generation, where the model learns to balance exploration (eliciting detailed preferences) and exploitation (offering confident,




Fig. 2. A single PEARL dialogue is annotated with structural tokens to train two separate simulators: *(left)* masks assistant turns and keeps user turns to train **UserSim**; *(right)* masks user turns and keeps assistant turns to train **RecSim**.

informed recommendations). Crucially, the recommender has no access to target items, compelling it to build genuine understanding and match user needs organically rather than guiding conversations toward predefined outcomes. Because our user simulator accepts any recommendation that satisfies its attribute criteria, conversations naturally conclude when a suitable match emerges. This avoids forced extensions seen in systems that must reach a specific target item. Instead, interactions end organically when users express satisfaction, resulting in more realistic dialogue patterns. The recommender learns to detect these signals and adapt its strategy, shaping nuanced, feedback-driven behaviors rather than following rigid scripts.

3.5 Independent Training for Specialized Behavior

Our novel training strategy is composed of two key components designed to foster specialized agent behavior. First, we enforce a structured output format to make conversational actions explicit. Second, we employ a role-specific loss masking technique to ensure each simulator learns exclusively from its own turns.

Structured Action Generation To enforce role-specific behavior, we structure the entire model output to separate the *intended action* from the *natural language response*. Each generated turn adheres to a strict format where an `<action>` block contains a specific command: `<recommend>`, `<inquire>`, `<greeting>`, `<disclose-goal>`, `<feedback>`, and `<accept>`. Afterwards, a `<response>` block provides the user-facing text, including the movie title, where applicable. These structural tokens are added to the tokenizer before training. This design provides two key advantages:

1. **Structured Reasoning:** It imposes an explicit response template, teaching the model to first commit to an unambiguous action before generating its utterance. For example, when providing a recommendation, the model can respond with:
`<action><recommend></action><response>You may enjoy <movie_title>Good Will Hunting (1997)</movie_title>.`
2. **Simplified Analysis:** The structured output format enables reliable post-simulation parsing, allowing us to programmatically extract and analyze the sequence of actions and recommended items. Crucially, the user simulator is trained to output `<accept>` when it encounters a satisfactory recommendation. This provides a natural and concrete stopping signal, even if the original oracle item is never recommended. In contrast, oracle-based methods can only terminate once the ground-truth item is recommended, often leading to either excessively long conversations or manual intervention where the ground truth is injected after a fixed number of turns — an unrealistic shortcut. Our reference-free design avoids this rigidity: conversations can conclude flexibly as soon as user satisfaction is reached. While the possibility of indefinitely extended dialogues still exists in principle, in practice this approach greatly reduces the need for manual intervention and more closely reflects the open-ended nature of real human–recommender interactions.

Role-Specific Loss Masking Training two separate simulators on a single conversational stream presents a challenge: each model must leverage the full dialogue for context while only learning to produce its own role-specific responses. We address this with our key innovation: a token masking strategy during loss computation. We calculate cross-entropy loss only over tokens corresponding to the simulator’s active role:

$$L_{\text{masked}} = - \sum_{t=1}^T M_t \cdot \log P(y_t | y_{<t}, x) \quad (1)$$

where the mask M_t is defined as:

$$M_t = \begin{cases} 1 & \text{if } t \in T_{\text{simulator}} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

This allows each model to see the full conversation for context while learning to generate only role-appropriate responses. This training approach avoids the role-swapping issues seen when a single model plays multiple roles. Instead, each simulator develops distinct conversational strategies: the user learns to express preferences naturally, while the recommender learns to generate recommendations and persuasive explanations. The result is two specialized models that understand dialogue dynamics but exhibit clearly differentiated personas.

3.6 Emergent Realistic Dynamics

When our independently trained simulators interact, their distinct objectives and perspectives produce authentic conversational dynamics that single-model setups fail to capture. The user simulator enters with clear yet flexible goals, expressing preferences naturally and staying open to unexpected but fitting suggestions. Meanwhile, the recommender must genuinely interpret these preferences, asking clarifying questions and generating recommendations based on the evolving dialogue, not preloaded answers. This independence is key to realism: the natural tension between user uncertainty and recommender confidence emerges organically from their interaction.

4 Experimental Setup

4.1 Experimental Details

We perform preprocessing steps to filter PEARL [12]. Specifically, we removed turns with empty dialogues and removed conversations with movies that were not part of the IMDB dataset used to generate PEARL. We filtered 7.65% of the original dataset, leaving our final dataset with 52.9k conversations. We fine-tune our both simulators using Llama 3.1 8B [5] and Qwen 3 8B [24]. Following Kong et al. [13], we use the base model versions rather than the instruct version to avoid the issue of pre-existing system alignment.

We perform full fine-tuning on both models for 2 epochs on 4 Nvidia A100 GPUs with a batch size of 4, gradient accumulation steps of 4, and a learning rate of $2e-4$. We evaluate on the filtered test set of the PEARL dataset, containing 2,069 conversations with seen user personas but unseen target items. For full implementation details, see: <https://github.com/jeromeras70/interplay-sim>.

4.2 Baselines

Although previous works use closed-sourced models GPT3.5 or GPT4 to simulate conversations, we rely only on open-source baselines for reproducibility. Namely, we use the instruction-tuned versions of Llama3.1 70B [5] and Qwen3 32B [24] as zero-shot user and recommender simulators. Furthermore, we fine-tune UniCRS [26], a popular CRS model, on the PEARL dataset as another baseline.

For our zero-shot baselines, we construct prompts that mirror the input to our fine-tuned models. For the user simulator role, the prompt includes the user’s persona (general preferences, movie history with reviews) and the target attributes. For the recommender simulator role, the prompt includes the user’s persona and the full conversation history. This setup ensures a fair comparison by providing the baselines with the same contextual information as our models.

4.3 Evaluation Metrics

We assess our simulators with automatic metrics on two fronts: (i) *multi-turn simulation quality* and (ii) *single-turn response quality*. Human evaluations are described in Section 5.3.

User Simulation Quality Each test dialogue has a known ground-truth target item. As the recommender proposes items, the simulator either accepts or rejects them. We summarize outcomes with three categories:

- **Success Rate (SR):** The simulator accepts the ground-truth item.
- **Early Termination (ET):** The simulator accepts some *non-ground-truth* item *before* the ground-truth is even proposed.
- **Failure Rate (FR):** The simulator rejects all recommendations, including the ground-truth item.

Higher SR and lower FR indicate that the simulator reliably follows user goals and accepts valid recommendations. A moderate ET rate can also be desirable: it reflects that the simulator occasionally accepts alternative items instead of always defaulting to the ground-truth recommendation, capturing a more diverse and realistic range of user preferences. By definition, these three rates are mutually exclusive and together account for 100% of dialogues.

Single-Turn Response Quality To evaluate the quality of individual user responses, we report:

- **BERTScore:** We measure the semantic similarity between the generated user responses and the ground-truth responses at each turn to evaluate contextual relevance (higher is better) [29].
- **Dist-4:** Measures lexical diversity as the fraction of unique 4-grams in the generated responses, computed over the whole test set (higher is better).

Recommender Performance In addition to Dist-4, we report the following metrics for the recommender:

- **Recall@1:** We check if the target item in a given turn is correct because our model is trained to only output a single recommendation in a single turn (higher is better).
- **Match Score:** To better evaluate recommendation quality in a reference-free context, we introduce Match Score. This metric quantifies the average similarity between the recommended items and the ground-truth movie. The score is the cosine similarity between movie embeddings from a model trained on the MovieLens-32M dataset [6]. This model uses TinyBERT [8] to encode semantic features (e.g., title, tags, genres) and learns interaction relevance via BPR loss [21] to capture user preferences (higher is better).

5 Results and Discussions

5.1 Realistic User Simulation without Oracle Knowledge

Multi-turn user simulation We evaluate whether our reference-free user simulator can appropriately accept and reject recommendations in a multi-turn setting by replaying conversations from the PEARL test set. Specifically, the user simulator walks through the recommender’s responses at each turn and determines

Table 1. Evaluations for User Simulators.

Model	Multi-Turn			Single-Turn		
	SR	ET	FR	BertScore	Dist-4	# Words
UserSim-Llama 8B	0.9265	0.0498	0.0237	0.9364	0.3023	31.2219
UserSim-Qwen 8B	0.9531	0.0232	0.0237	0.9373	0.3093	30.7637
Llama3.1 70B	0.3615	0.0033	0.6351	0.8870	0.2786	43.4067
Qwen3 32B	0.7772	0.0097	0.2131	0.8921	0.3409	31.6797

whether to accept or reject the item. This setup directly tests the simulator’s ability to handle ongoing dialogue without oracle knowledge. As shown in Table 1, our reference-free user simulator achieves high success rates and effectively addresses the oracle problem.

Our fine-tuned models achieve success rates of 93-95%, substantially outperforming larger zero-shot LLMs including Llama3.1 70B (36% SR) and Qwen3 32B (77% SR). Critically, these results show that our simulators can recognize and accept appropriate recommendations without knowing the target movie title in advance, validating that reference-free operation preserves realistic user behavior while eliminating artificial constraints.

Single-turn response quality We also evaluate the user simulator’s single-turn response quality by generating a user reply at each turn t using the chat history up to that point ($h = \{u_1, a_1, \dots, a_{t-1}\}$). Table 1 shows our models achieve superior BertScore compared to larger baselines while maintaining good diversity. This confirms that removing target knowledge does not compromise the quality of user responses – our simulators generate contextually appropriate utterances that match human references without requiring predetermined outcomes.

5.2 High-Quality Recommendations without Oracle Knowledge

To determine whether the model can provide high-quality recommendations without oracle knowledge, we generate recommendations from our model on the turn when the ground truth is given. To measure the recommendation performance of our model, we report the Match Score between the recommended items and the ground truth movie per dialogue round.


Table 2 provides strong evidence that specialized training enables effective recommendation without oracle knowledge. Our RecSim-Qwen8B achieves the highest performance with Recall@1 of 0.0217 and Match Score of 0.9333, significantly outperforming the larger Qwen3-32B baseline and exceeding the modular UniCRS [26] system. This substantial performance gap demonstrates that fine-tuning smaller models for specific conversational recommendation tasks is more effective than relying on larger general-purpose LLMs with predetermined targets.

Additionally, Table 2 shows our models generate more diverse responses than larger baselines while maintaining appropriate lengths. This indicates that our

Table 2. Evaluations of recommendation performance.

Model	Recall@1	Match Score	Dist-4 #	Words
RecSim-Llama8B	0.0193	0.9323	0.5243	37.42
RecSim-Qwen8B	0.0217	0.9333	0.5358	37.30
Qwen3 32B	0.0014	0.4745	0.4630	42.21
UniCRS	0.0119	0.9129	0.2147	29.69

approach produces contextually rich, varied recommendations without the repetitive patterns common in oracle-based systems.

**Fig. 3.** Human evaluation win ratio comparing reference-free dialogues to PEARL dataset. (* indicates p-value < 0.05)

5.3 Realistic Dialogue via Independent Simulation

To assess whether our framework generates realistic conversations, we conducted a human evaluation with 10 users across 50 dialogues pairs comparing our reference-free dialogues against PEARL dataset conversations using six criteria from Kim et al. [12]: user control, expertise, specificity of preferences, relevance, conversational flow, and consistency. Specifically, human judges were shown a head-to-head comparison between two dialogues generated using the same user profile and were instructed to select which dialogue performed better.

Figure 3 reveals that our reference-free approach achieves a quality comparable to PEARL’s reference-dependent conversations. Evaluators found that our reference-free dialogues exhibited significantly more user control, though they did not match the naturalness of PEARL. This gap is partly attributable to how the models are trained: because our simulators are trained on PEARL’s LLM-generated text, they may inherit its stylistic patterns, so perceived naturalness likely reflects the training data’s style rather than a limitation of our reference-free design. PEARL was also generated using GPT-3.5, which has substantially

Table 3. A example of the head-to-head comparison of our **Interplay** simulator (Dialogue A) versus a **PEARL** dialogue (Dialogue B). The table highlights the key turns in the conversations. Notably, **Interplay** can make high-quality recommendations without any predefined movies or target items to recommend.

Dialogue A (Interplay)	Dialogue B (PEARL)
<i>Initial user preference is established...</i>	
Turn 3 - User: I really enjoy watching movies with strong character development and intense chase/fight scenes. I also appreciate well-written scripts and great performances...	Turn 3 - User: I'm looking for something with action and adventure. I enjoy movies with thrilling chase and fight scenes, as well as strong performances...
<i>...after several exchanges and rejected recommendations...</i>	
Turn 7 - User: I appreciate the suggestion, but I found <i>Superman Returns</i> to be lacking in intensity... the pacing felt off... I'm hoping for something with a little more edge and excitement.	Turn 7 - User: I appreciate the recommendation, but I'm not a big fan of <i>Live and Let Die</i> ... I found the villains to be a bit too over-the-top... I tend to prefer movies with a more serious tone...
Turn 8 - Rec: How about giving <i>The Bourne Identity (2002)</i> a try? It has everything you're looking for - strong character development, intense chase and fight scenes, a well-written script, and excellent acting performances.	Turn 8 - Rec: How about watching <i>Quantum of Solace (2008)</i> ? It's a James Bond film that has thrilling chase and fight scenes, as well as incredibly well-written plot.
Turn 9 - User: That sounds perfect! I haven't watched <i>The Bourne Identity</i> yet, so I'll definitely check it out. Thanks for the recommendation!	Turn 9 - User: That sounds perfect! I haven't seen <i>Quantum of Solace</i> before, so I'll definitely give it a watch. Thanks for the recommendation!

more parameters than our 8B model, further contributing to its higher naturalness ratings. Nevertheless, our user study shows that eliminating oracle knowledge does not compromise dialogue authenticity: independent simulators produce realistic interactions through genuine exploration rather than scripted convergence on predetermined targets. Moreover, our approach achieves this quality with smaller models, significantly reducing computational costs while maintaining conversational realism.

6 Limitations

Domain Specificity Our empirical evaluation is confined to the movie recommendation domain, chosen due to the availability of large-scale, high-quality public datasets like **PEARL** [12]. While the **Interplay** framework is designed to be domain-agnostic, its effectiveness in other domains such as e-commerce, music, or travel has not yet been validated. Future work should focus on applying and evaluating our independent simulator approach in these different contexts, which may present unique challenges like different conversational dynamics.

Constrained Conversational Scope The current simulation operates within a structured, task-oriented dialogue framework where conversational turns are mapped to explicit actions like `<recommend>` or `<feedback>`. This approach is beneficial for control and analysis but does not capture the full spectrum of real-world interactions. Human conversations often involve open-domain chit-chat, multi-intent utterances (e.g., asking a question and providing feedback in the same turn), and more fluid topic shifts. A key direction for future research is to enhance the simulators to handle these less structured, mixed-initiative conversations.

Limitation on Recommendation Evaluation Although MatchScore is a useful automatic proxy for recommendation quality, it has not yet been validated as a reliable metric that aligns with real user preferences. Future work should include user studies to evaluate whether automatic metrics such as MatchScore correlates with human judgments, and more broadly evaluate whether the reference-free simulation results reflect real-world user behavior.

7 Conclusion

We tackle the oracle problem in conversational recommender system simulation, where access to target items creates unrealistic, constrained interactions. Our reference-free framework trains independent user and recommender simulators without predetermined outcomes, enabling more natural conversations and flexible recommendation behavior. By focusing on target attributes rather than items, user simulators accept multiple valid suggestions, while recommender simulators develop genuine preference understanding without relying on oracle knowledge.

Our approach outperforms larger models across key metrics: user simulators achieve 93–95% success rates without target leakage, and our 8B recommenders surpass 32B baselines by $15\times$ in accuracy while maintaining high diversity. Human evaluations confirm our dialogues match reference-dependent quality, but with far lower computational cost. This establishes a scalable, realistic simulation paradigm for CRS and paves the way for more advanced multi-agent systems in conversational AI.

8 Acknowledgements

This project was supported using computational resources from the Nvidia Academic Grant Program.

9 Disclosure of Interests

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

Bibliography

- [1] Zhendong Chu, Nan Wang, and Hongning Wang. Multi-objective intrinsic reward learning for conversational recommender systems. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, Red Hook, NY, USA, 2023. Curran Associates Inc.
- [2] Jacob Eisenstein, Reza Aghajani, Adam Fisch, Dheeru Dua, Fantine Huot, Mirella Lapata, Vicky Zayats, and Jonathan Berant. Don't lie to your friends: Learning what you know from collaborative self-play, 2025. URL <https://arxiv.org/abs/2503.14481>.
- [3] Jiabao Fang, Shen Gao, Pengjie Ren, Xiuying Chen, Suzan Verberne, and Zhaochun Ren. A multi-agent conversational recommender system, 2024. URL <https://arxiv.org/abs/2402.01135>.
- [4] Yunfan Gao, Tao Sheng, Youlin Xiang, Yun Xiong, Haofen Wang, and Jiawei Zhang. Chat-rec: Towards interactive and explainable llms-augmented recommender system. *arXiv preprint arXiv:2303.14524*, 2023.
- [5] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.
- [6] F. Maxwell Harper and Joseph A. Konstan. The movielens datasets: History and context. *ACM Trans. Interact. Intell. Syst.*, 5(4), 2015. <https://doi.org/10.1145/2827872>. URL <https://doi.org/10.1145/2827872>.
- [7] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. A survey on conversational recommender systems. *ACM Comput. Surv.*, 54(5), May 2021. ISSN 0360-0300. <https://doi.org/10.1145/3453154>. URL <https://doi.org/10.1145/3453154>.
- [8] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling bert for natural language understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, 2020. <https://doi.org/10.18653/v1/2020.findings-emnlp.372>. URL <https://doi.org/10.18653/v1/2020.findings-emnlp.372>.
- [9] Hideaki Joko and Faegheh Hasibi. Face: A fine-grained reference free evaluator for conversational recommender systems, 2025. URL <https://arxiv.org/abs/2506.00314>.
- [10] Hideaki Joko, Shubham Chatterjee, Andrew Ramsay, Arjen P. de Vries, Jeff Dalton, and Faegheh Hasibi. Doing personal laps: Llm-augmented dialogue construction for personalized multi-session conversational search. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '24, page 796–806, New York, NY, USA, 2024. Association for Computing Machinery. ISBN

9798400704314. <https://doi.org/10.1145/3626772.3657815>. URL <https://doi.org/10.1145/3626772.3657815>.
- [11] Dongyeop Kang, Anusha Balakrishnan, Pararth Shah, Paul Crook, Y-Lan Boureau, and Jason Weston. Recommendation as a communication game: Self-supervised bot-play for goal-oriented dialogue. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1951–1961, Hong Kong, China, November 2019. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1203>. URL <https://aclanthology.org/D19-1203/>.
 - [12] Minjin Kim, Minju Kim, Hana Kim, Beong-woo Kwak, SeongKu Kang, Youngjae Yu, Jinyoung Yeo, and Dongha Lee. Pearl: A review-driven persona-knowledge grounded conversational recommendation dataset. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Findings of the Association for Computational Linguistics: ACL 2024*, pages 1105–1120, Bangkok, Thailand, August 2024. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.findings-acl.65>. URL <https://aclanthology.org/2024.findings-acl.65/>.
 - [13] Chuyi Kong, Yaxin Fan, Xiang Wan, Feng Jiang, and Benyou Wang. PlatoLM: Teaching LLMs in Multi-Round Dialogue via a User Simulator. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7841–7863, Bangkok, Thailand, 2024. Association for Computational Linguistics. URL <https://aclanthology.org/2024.acl-long.424>.
 - [14] Heejin Kook, Junyoung Kim, Seongmin Park, and Jongwuk Lee. Empowering retrieval-based conversational recommendation with contrasting user preferences. In Luis Chiruzzo, Alan Ritter, and Lu Wang, editors, *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7692–7707, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN 979-8-89176-189-6. URL <https://aclanthology.org/2025.naacl-long.392/>.
 - [15] Chuang Li, Hengchang Hu, Yan Zhang, Min-Yen Kan, and Haizhou Li. A conversation is worth a thousand recommendations: A survey of holistic conversational recommender systems. In *Proceedings of KaRS@RecSys 2023*. CEUR-WS.org, 2023. URL <https://ceur-ws.org/Vol-3560/long1.pdf>.
 - [16] Raymond Li, Samira Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. Towards deep conversational recommendations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, page 9748–9758, Red Hook, NY, USA, 2018. Curran Associates Inc.
 - [17] Tingting Liang, Chenxin Jin, Lingzhi Wang, Wenqi Fan, Congying Xia, Kai Chen, and Yuyu Yin. LLM-REDIAL: A large-scale dataset for conversational recommender systems created from user behaviors with LLMs.

- In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Findings of the Association for Computational Linguistics: ACL 2024*, pages 8926–8939, Bangkok, Thailand, August 2024. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.findings-acl.529>. URL <https://aclanthology.org/2024.findings-acl.529/>.
- [18] Ahtsham Manzoor, Samuel C. Ziegler, Klaus Maria. Pirker Garcia, and Dietmar Jannach. 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, page 267–272, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704338. <https://doi.org/10.1145/3627043.3659574>. URL <https://doi.org/10.1145/3627043.3659574>.
 - [19] Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. OpenDi-alkG: Explainable conversational reasoning with attention-based walks over knowledge graphs. In Anna Korhonen, David Traum, and Lluís Màrquez, editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 845–854, Florence, Italy, July 2019. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1081>. URL <https://aclanthology.org/P19-1081/>.
 - [20] Zhangchi Qiu, Linhao Luo, Shirui Pan, and Alan Wee-Chung Liew. Unveiling user preferences: A knowledge graph and llm-driven approach for conversational recommendation, 2024. URL <https://arxiv.org/abs/2411.14459>.
 - [21] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2009. URL https://www.auai.org/uai2009/papers/UAI2009_0139_48141db02b9f0b02bc7158819ebfa2c7.pdf.
 - [22] Yueming Sun and Yi Zhang. Conversational recommender system. In *The 41st international acm sigir conference on research & development in information retrieval*, pages 235–244, 2018.
 - [23] Rohan Surana, Junda Wu, Zhouhang Xie, Yu Xia, Harald Steck, Dawen Liang, Nathan Kallus, and Julian McAuley. From reviews to dialogues: Active synthesis for zero-shot llm-based conversational recommender system, 2025. URL <https://arxiv.org/abs/2504.15476>.
 - [24] Qwen Team. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.
 - [25] Tao Tu, Anil Palepu, Mike Schaeckermann, Khaled Saab, Jan Freyberg, Ryutaro Tanno, Amy Wang, Brenna Li, Mohamed Amin, Nenad Toma-sev, Shekoofeh Azizi, Karan Singhal, Yong Cheng, Le Hou, Albert Web-son, Kavita Kulkarni, S Sara Mahdavi, Christopher Sementur, Juraj Got-tweis, Joelle Barral, Katherine Chou, Greg S Corrado, Yossi Matias, Alan Karthikesalingam, and Vivek Natarajan. Towards conversational diagnostic ai, 2024. URL <https://arxiv.org/abs/2401.05654>.
 - [26] Xiaolei Wang, Kun Zhou, Ji-Rong Wen, and Wayne Xin Zhao. Towards unified conversational recommender systems via knowledge-enhanced prompt

- learning. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pages 1929–1937, 2022.
- [27] Xiaolei Wang, Xinyu Tang, Xin Zhao, Jingyuan Wang, and Ji-Rong Wen. Rethinking the evaluation for conversational recommendation in the era of large language models. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10052–10065, Singapore, December 2023. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.emnlp-main.621>. URL <https://aclanthology.org/2023.emnlp-main.621/>.
 - [28] Shuo Zhang and Krisztian Balog. Evaluating conversational recommender systems via user simulation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '20, page 1512–1520, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379984. <https://doi.org/10.1145/3394486.3403202>. URL <https://doi.org/10.1145/3394486.3403202>.
 - [29] Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SkeHuCVFDr>.
 - [30] Kun Zhou, Wayne Xin Zhao, Shuqing Bian, Yuanhang Zhou, Ji-Rong Wen, and Jingsong Yu. 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, page 1006–1014, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379984. <https://doi.org/10.1145/3394486.3403143>. URL <https://doi.org/10.1145/3394486.3403143>.
 - [31] Lixi Zhu, Xiaowen Huang, and Jitao Sang. How reliable is your simulator? analysis on the limitations of current llm-based user simulators for conversational recommendation. In *Companion Proceedings of the ACM Web Conference 2024*, WWW '24, page 1726–1732, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400701726. <https://doi.org/10.1145/3589335.3651955>. URL <https://doi.org/10.1145/3589335.3651955>.
 - [32] Lixi Zhu, Xiaowen Huang, and Jitao Sang. A llm-based controllable, scalable, human-involved user simulator framework for conversational recommender systems. In *Proceedings of the ACM on Web Conference 2025*, WWW '25, page 4653–4661, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400712746. <https://doi.org/10.1145/3696410.3714858>. URL <https://doi.org/10.1145/3696410.3714858>.