



From Recommendation to Exploration: A Human-Centered Review of Conversational Recommender Systems

Chi-hsuan Lin

Abstract

Conversational Recommender Systems (CRSs) have emerged as an important research area that combines recommendation algorithms with natural language interaction to provide more personalized and engaging user experiences. While current studies mostly focus on improving accuracy, dialogue management, and model performance, less attention has been given to user experience. This literature review examines the shift from system-centered research toward human-centered CRS design by synthesizing studies on explainability, trust, discoverability, preference elicitation, and exploratory information seeking. It compares different ways that help the user who is not sure what they are looking for to explore information instead of providing the most accurate recommendations. This review also discusses the current method for evaluating users' experience and some improvements that could be made to it. Overall, this review argues that future research ought to not only be done towards improving recommendation accuracy but also on user experience by helping them to understand, explore, and make better decisions.

Keywords: Conversational Recommender Systems, Natural language interaction, Human-centered CRS, Explainability, Exploration, User experience

1 Introduction

1.1 Background and Context

With the increase in online information, finding relevant information has become harder for users. Recommendation systems help solve this problem by suggesting products, movies, music, or other content based on a user's interests and previous interactions [7, 12, 20]. However, traditional recommendation systems mainly rely on historical data and often assume that users already know what they are looking for. This can make it difficult for users to discover new interests or explore unfamiliar options [7, 14].

Conversational Recommender Systems (CRSs) extend traditional recommendation systems by having a dialogue with the user to take in current preferences from the user, allowing the user to interact with the system through natural language conversations [7, 10, 14]. Instead of only providing recommendations, these systems can ask follow-up questions, explain their suggestions, and adjust recommendations as users share more information [10]. As large language models (LLMs) become more advanced, CRSs also become more capable of engaging and interacting with the user [20, 25].

1.2 Research Motivation

Although many studies have focused on improving recommendation accuracy and conversation strategies, particularly through preference elicitation, dialogue management, and strategic planning [4, 7, 14, 25], less attention has been given to how users experience these interactions. Recent research has begun to explore how explainability, trust, discoverability, and preference elicitation can help users better explore the possibilities and make decisions [6, 9, 11, 20]. These human-centred aspects are important to help users discover all kinds of possibilities instead of simply receiving the most accurate recommendation. To better understand the shift from recommendation performance toward broader user support, this literature review explores the following research question: ***How can Conversational Recommender Systems support users in developing, exploring, and understanding their preferences?***

This literature review examines how recent research is shifting from improving recommendation performance to improving the overall user experience of Conversational Recommender Systems. It focuses on four areas that emerged from the literature as different ways CRS can support users throughout the recommendation process: discoverability and information exploration, which help users find and explore options, and explainability and trust, which help users understand and evaluate recommendations. Together, these factors can help conversational systems better support users throughout the recommendation process.

2 Methods

2.1 Search Strategy

To provide an overview of how CRSs have improved and the future improvements towards user experiences, we identified, reviewed, and selected relevant literature using PRISMA guidelines. To ensure the studies are applicable, we considered 2 factors: (1) the number of times it was cited and (2) time — after 2018. We used Google Scholar and ACM Digital Library to find research related to Conversational Recommender Systems with search terms such as “conversational recommendation”, “conversational recommender system,” “explainable conversational recommendation,” and “user experience.”

2.2 Inclusion Criteria

The following are the criteria to identify whether a study is related to Conversational Recommender Systems:

- The paper was published after 2018.
- The paper is written in English.

- The paper focuses on Conversational Recommender Systems (CRSs) or closely related conversational recommendation technologies.
- The paper was peer-reviewed.
- The paper discusses at least one aspect of CRS design, performance, interaction, or user experience that could contribute to understanding how CRSs support users beyond recommendation accuracy.

The initial search contains 91 studies. We evaluated these studies based on the title and abstract; 85 studies remained. After reading and reviewing these studies fully, we selected 23 relevant studies that meet the criteria for analysis (see Figure 1).

The papers reviewed ($n = 23$) were published between 2018 and 2026, with an increasing number in 2024 (see Figure 2). Most of the papers were published in the ACM Digital Library.




Figure 1: PRISMA Flow Chart [3]




Figure 2: Number of Papers by Year

2.3 Analysis Method

The selected papers were grouped according to the research themes, including adapting to users' preferences, supporting exploration and discoverability, improving understanding by

explanation, and creating a user-centered interaction. The similarities and differences between studies were then compared to identify the trends and research gaps.

3 Results

In this study, we suggest that supporting users in conversational recommender systems (CRSs) involves more than generating accurate suggestions. While recommendation accuracy remains an important measure of CRSs' effectiveness, recent research has increasingly examined how systems can support users throughout the recommendation process. Across this literature, four main forms of user support were shown: understanding and adapting to users' preferences, facilitating exploration and discoverability, helping users understand recommendations through explanations, and creating human-centered interactions. These dimensions highlight how CRSs can move beyond simply predicting users' preferences towards supporting users in exploring possibilities, making decisions, and engaging in meaningful conversations.

3.1 Supporting Users Through Adaptive Preference Understanding

Traditional recommender systems primarily rely on historical user data, provide limited interaction with users, and provide a one-shot recommendation [1, 7]. To tackle these limitations, Conversational Recommender Systems (CRSs) are needed. These types of systems provide real-time interactions and negotiated dialogues with users, enabling recommendations to take in current preferences from users and be updated throughout the interaction [10, 14]. This interactive process provides an opportunity for CRSs to move beyond sticking with a fixed user's preference and instead continuously adapt to the user's needs as the conversation develops.

Recent research has therefore focused on making preference understanding more adaptive and responsive. Rather than relying on a single preference or treating each user response independently, researchers have explored methods that allow CRS to plan conversations, incorporate feedback, and continuously update their understanding of users. These approaches aim to address situations where the users have multiple, incomplete, changing, or unreliable preferences.

One recent approach is SAPIENT, a framework for conversational recommender systems to help the AI agent think in the long term and train it to perform better as it goes [4]. Rather than reacting to each user's messages, SAPIENT considers different actions that can be taken, such as asking additional questions, providing explanations, or making recommendations, to select the best actions to maximize user satisfaction. This type of approach allows the system to determine which information is necessary and reduce redundant questions. By planning across multiple turns, SAPIENT produces more efficient dialogues and improves recommendation quality, conversation success rate, and overall user engagement.

While SAPIENT focuses on planning conversational actions, Reinforced Strategy Optimization (RSO) focuses on adapting conversational strategies to different users and conversation states [25]. It splits the user input into two parts: (1) Macro-level Strategy Planning, which determines the overall conversational strategy and selects which type of action to take (e.g., recommend, explain, encourage) based on the conversation state and user feedback; and (2) Micro-level Strategy Adaptation, which operationalizes the chosen strategy by adapting it to the user's context. With this strategy, it lowers the difficulty and complexity of the system and enables more personalized strategy learning. Overall, the study demonstrates that strategy-level adaptation can improve conversational quality and recommendation performance, while allowing the systems to respond more flexibly to users [25].

Other researchers have focused on improving the CRS model for user preferences during the conversation. The Multi-Interest Multi-Round Conversational Recommendation System with Fuzzy Feedback Based User Simulator recognized that users often have multiple preferences that cannot be represented by a single fixed profile [19]. Its use of fuzzy feedback allows the system to represent uncertainty instead of forcing a user's preferences into precise rankings.

Uncertainty also arises when user feedback is unreliable. Online Learning and Detecting Corrupted Users for Conversational Recommendation Systems address the underexplored challenges of misleading or corrupted feedback from users in CRSs by enabling the system to detect it during the learning process [2]. By filtering the unreliable feedback, the system can improve its accuracy in modeling user preferences and recommendations.

Another issue in conversational recommender systems is confounding bias, where previous system interactions influence the feedback used to predict user preferences. To address this issue, researchers came up with the User-Regulation Deconfounded Conversational Recommender System with Bandit Feedback, which uses bandit learning techniques to separate genuine user preferences from behavioral biases [23]. By reducing the influence of confounding factors, the model is able to produce recommendations that more accurately reflect the user's actual preferences. Together, these studies demonstrate that adaptive preference understanding involves more than collecting additional user information. Effective CRS must also determine how preferences change, how reliable feedback is, and how previous interactions may influence the interpretation of user behavior.

Taken together, these studies show that CRS can support users by treating ongoing preference understanding as an ongoing process rather than a one-time prediction. Planning conversational strategies, adapting to user feedback, and accounting for unreliable preferences allow systems to respond to users who are still uncertain about their needs. However, understanding users more effectively does not necessarily mean that users know what they want in the first place.

When users have uncertain preferences or are interested in discovering unfamiliar possibilities, CRS can play a different role: helping them explore

3.2 Supporting Exploration and Discoverability

Understanding users' existing preferences is only one way that conversational recommender systems can support them. Users do not always enter a conversational recommender with clearly defined preferences or options. In these situations, the role of CRS can extend beyond predicting what the user will likely prefer to help them discover relevant possibilities and formulate what they are looking for. Recent research has therefore examined how conversational systems can improve discoverability, proactively introduce relevant options, and encourage users to explore interests beyond their existing knowledge or historical preferences.

One challenge in conversational systems is that users may not be aware of all the capabilities or skills available to them, limiting their ability to benefit from the conversational assistant. Kalirai et al. [12] found that users typically interact with only a small number of available skills because they do not know what functions the assistant offers. To solve this issue, the authors proposed faceted skill recommendations, which suggest relevant tools based on the user's interests. Rather than requiring users to know what capabilities to ask for, the system introduces potentially useful functions during the interaction. This demonstrates how conversational recommender systems can not only provide suggestions on relevant items, but also help users discover what the system can help them accomplish.

Similar ideas are explored by Liao et al. [5], who examined discoverability in enterprise conversational assistants. Using real-world data from the AI Assistant for Adobe Experience Platform, the researchers discovered how users generally use a small fraction of features. To address this problem, the authors propose skill question suggestions, which use previous interactions to generate questions that may be useful to the user. These suggestions can help users discover additional system capabilities and encourage users to explore topics they might not have considered. Compared with faceted skill recommendation, which primarily focuses on exposing available system functions, proactive question suggestions guide users toward asking better questions, thereby formulating new directions for the conversation.

Exploration can also extend beyond discovering system capabilities to discovering new personal interests. Manveer et al. [10] stated that future recommender systems should support and nurture user self-actualization rather than simply use a fixed preference inferred from historical interactions. Interactive conversational recommendation could help the agent take in real-time information to ensure the recommendations are suitable for the users' current preferences. This system also encourages users to discover new interests, pursue future goals, and explore identities that may not be reflected in historical interaction data. From this perspective, CRS

does not merely infer what users already want; it can participate in the process through which users explore and develop new preferences.

Together, these studies demonstrate that exploration in CRS can occur at multiple levels. Faceted skill recommendations and skill question suggestions improve the discoverability of system capabilities and help users formulate new directions for conversation, while work on self-actualization extends exploration to users' own developing interests and goals. These approaches shift the role of conversational recommendation from predicting a fixed preference to supporting users as they discover new possibilities. However, greater exploration also introduces a need for users to understand and evaluate unfamiliar recommendations, making explanation an important component of the exploratory experience.

3.3 Supporting Understanding Through Explanations

Exploration does not end when a CRS presents an unfamiliar option. Users also need to understand why an option may be relevant and determine whether it aligns with their interests. Explanations can therefore serve not only to justify a recommendation, but also to provide information that helps users evaluate unfamiliar possibilities. Recent research has examined how the content and quality of explanations influence users' understanding, trust, and willingness to explore recommendations.

Explanation Driving Exploration argues that conversational recommender systems should explain recommendations to users in a way that helps them explore unfamiliar content, not just accept recommendations blindly [11]. By explaining a decision or option the CRS recommends to the user, the system becomes more trustworthy to the user, and it helps the user to understand the recommendation better. Through user studies, the author identified four dimensions of information that people naturally seek when exploring unfamiliar items: essence, experience, exchange, and entwinement. While humans typically combine multiple dimensions, many existing systems only provide one type of explanation. Providing explanations that include multiple information dimensions can therefore improve users' understanding of recommendations and encourage deeper exploration. These findings suggest that explanation quality plays not only a technical role in a recommendation system, but also an important part of how users make sense of and evaluate recommendations.

The importance of explanations is also reflected in research on user-centered evaluation. CRS-Que examines how explainability and transparency contribute to users' experience with CRSs [9]. Its findings suggest that explanations can increase users' sense of control by giving them greater insight into the recommendation process. This indicates that explanations can serve as a broader function than providing reasons for recommendations: they can help users

feel more informed and involved in the recommendation process. In this way, explainability connects the recommendation itself with the user's ability to understand and evaluate it.

Other explainable CRS research similarly demonstrates the relationship between explanation quality and trust. Explainable conversational recommendation approaches have been developed to generate explanations alongside recommendations, allowing users to better understand the reasons for the recommended items. Research has shown that improving the quality and informativeness of these explanations can increase the explainability and trustworthiness of the CRS [22]. More recent work also highlights the importance of credibility, as explanations containing misinformation can undermine users' trust even when they are persuasive [18]. Therefore, simply providing an explanation won't necessarily improve user experience; explanations must also provide credible, understandable, and relevant information.

As a whole, these studies suggest that explanations can support users at several stages of the recommendation process. They can help users understand unfamiliar options, evaluate the effectiveness of recommendations, and develop greater trust during the interaction. Importantly, this means that explanation quality should not be evaluated solely according to whether a system can generate a plausible justification. From a user-centered perspective, explanations are valuable when they provide the information users need to understand and make decisions about recommendations. This extends the role of CRS beyond selecting relevant options toward helping users actively evaluate and make sense of those options.

3.4 Human-Centered Interaction

The previous sections show that CRSs can support users by adapting to their preferences, helping them explore different possibilities, and providing explanations for recommendations. However, these functions are all experienced through the interaction between the user and the system. A CRS can provide relevant recommendations but still be difficult to use if it does not understand the user, respond appropriately, or communicate naturally. Therefore, recent research has also examined CRS from a human-centered perspective, focusing on how users experience and interact with these systems.

One of the first comprehensive frameworks developed specifically for evaluating conversational recommender systems is CRS-Que [9]. Unlike traditional evaluation methods that primarily measure recommendation accuracy or dialogue success, CRS-Que evaluates the entire conversational experience from the user's perspective. The framework incorporates 6 major critical conversation constructs: CUI Understanding, CUI Adaptability, CUI Response Quality, CUI Attentiveness, CUI Rapport, and CUI Humanness. With these constructs, studies on user experience and user satisfaction can be conducted. They measure how naturally the system communicates with users. This was validated through user studies involving music exploration

and mobile phone purchase, which demonstrate good reliability in assessing user experience. The author also discovered that explainability increased users' control, and increasing the humanization level could increase users' trust. These findings suggest that conversational quality is closely connected to users' willingness to continue using a system.

While CRS-Que focuses primarily on users' experience during the conversation with the agent, CONCEPT considers both system- and user-centric factors to evaluate conversational recommender systems. CONCEPT is a tool that provides an overview of the pros and cons of CRS models and offers a reference guide to evaluate CRS [6]. It evaluates how technical system performance affects users' experience. This evaluation demonstrates that recommendation quality and users' experience are correlated. By considering both factors, CONCEPT provides a better and more accurate assessment for CRSs. As a result, researchers found the limitations in current CRS models and suggested that even "omnipotent" systems lack practical usability. With CONCEPT, a guideline or reference for evaluating CRSs has been provided.

Beyond developing evaluation frameworks, researchers have investigated the system qualities that enhance user experience. CRS-UX is a model to capture key qualities of conversational recommender systems from the users' perspective [8]. Using this model, researchers found that recommendation relevance, transparency, adaptability, trustworthiness, and conversational naturalness contribute to positive user experiences. With this information, they were able to improve the system's quality and evaluate the key factors in improving user experience. Moreover, researchers came up with the thought that speech-based conversational recommenders, such as intelligent voice assistants, are becoming prevalent. As a whole, these findings reinforce the idea that recommendation quality alone cannot fully explain users' satisfaction with conversational systems.

Several recent studies suggest that how conversational recommender systems recommend is equally important as what the systems recommend. Different conversational styles can influence users' experience. Research suggests that adapting conversational strategies to individual users can improve overall conversion quality and engagement [13]. Furthermore, another study found that different users significantly prefer different conversational systems to communicate. With this being said, researchers indicate that conversational recommender systems that can adapt according to the users' preferred style can further improve the overall experience [15]. More broadly, the design of a conversational agent's dialogue shapes user satisfaction. Its structure, format, and style affect users' trust, the agent's intelligence, and overall credibility [21]. Together, these studies show the importance of the style and design of the system. It is an essential component that strengthens users' satisfaction while using this system. Systems that can adapt to their users are key to improving conversational recommender systems.

Improvements in large language models (LLMs) have also contributed to the improvement of user experiences in conversational recommender systems. A study reported that LLM-based CRSs produce conversations that users described as more natural, flexible, and engaging compared to traditional CRSs. This system can clarify implicit needs, support unique explorations, and facilitate a deeper understanding of the topic the user wants to explore [24]. However, there are still ongoing challenges with this system, including occasional inaccuracies, inconsistent recommendations, and concerns regarding transparency. These findings suggest that although LLMs improved the quality of the conversation and user interactions, further improvements to this system are still needed to provide an overall more satisfying user experience.

Across these papers and studies, they all show that supporting users involves more than improving the recommendations themselves. Users also need a system that can understand their needs, adapt during the conversation, respond appropriately, and communicate in a way that feels natural and trustworthy. Research such as CRS-Que, CONCEPT, and the study of key CRS qualities provides different ways of evaluating these user-centered factors, while studies on conversational strategies and LLM-powered CRS show how interaction can be adapted and improved. These findings suggest that a human-centered approach is important for designing CRS supportive experiences throughout the conversation.

4 Discussion

The findings of this review suggest that the role of Conversational Recommender Systems is becoming broader than simply predicting what users may prefer. Although recommendation accuracy is still important, the literature suggests that users may need support at different stages of the recommendation process, such as developing their preferences, discovering unfamiliar options, understanding recommendations, and interacting with the system. These findings raise a broader question about what the main purpose of a CRS should be: should it not only predict what users are likely to prefer, or should it also help users explore and make decisions?

4.1 From Recommendation to Exploration

One important finding from this review is that users do not always know exactly what they want when they begin a conversation with a recommender system. Their preferences can change as they receive new information or discover options they had not considered before. Researchers argue that future recommender systems could support users not only with their current preferences but also with their future goals and personal development [10]. This means that a CRS cannot always treat user preferences as something that can simply be collected at the beginning of a conversation and used to produce one recommendation.

Instead, the conversation itself can become a part of the decision-making process. When a system adapts to changing preferences, introduces new possibilities, and provides information about recommendations, it can help users understand their own preferences as they interact with the system. This is different from a traditional recommender system, where the main goal is often to predict what the user will like based on existing data. Research on discoverability and exploration similarly shows that CRS can introduce system capabilities, suggest questions, and help users discover options that they may not have considered before [11, 12]. Explanations can also help users understand unfamiliar recommendations and decide whether they are relevant to their needs [11].

This also connects the different areas found in this review. Adaptive preferences understanding allows the system to respond when users' interests change. Exploration gives users the opportunity to consider options outside their initial request. Explanations can then help users understand and compare these options. These functions work together rather than separately, since users may need to discover an option before they can understand whether it is suitable for them.

Therefore, the role of conversational recommender systems should not simply be finding the most relevant recommendation. It could also act as a form of decision support by helping users explore their choices and make their own decisions. However, this raises another question: **how much should a CRS guide users without taking control away from them?** If a system becomes too proactive, it may limit the user's choices instead of supporting them. Future CRS research should therefore consider not only how systems can help users, but also how much control users should have over the interaction.

4.2 Rethinking How CRS Should Be Evaluated

The findings also show that recommendation accuracy does not tell us everything about whether a CRS is useful. Accuracy measures whether a system can provide relevant recommendations, but it does not tell us whether users understand those recommendations, feel that the system understands them, or find the conversation helpful. Human-centered CRS research has therefore considered factors beyond recommendation performance, including users' perceptions of the interaction [8, 9].

This is important because the purpose of a conversational system is different from that of a system that only produces recommendations. Users are actively involved in the interaction, which means that their experience can affect their thoughts on the usefulness of the system. For example, a system may recommend an item that matches the user's preferences but still provide a poor experience if it does not respond to their questions or adapt when their preferences change [13, 21].

The human-centered studies in this review suggest that these factors should be considered when evaluating CRS. Measures such as adaptability, understanding, response quality, transparency, trustworthiness, and conversational naturalness provide information that recommendation accuracy cannot capture [8, 9]. These findings suggest that two CRS with similar recommendation accuracy could still provide different experiences for users.

However, this does not mean that accuracy should be ignored. The issue is not whether accuracy or user experience is more important. Instead, the two types of evaluation can provide different information about the system. Accuracy can show whether a CRS is providing relevant recommendations, while user-centered measures can show whether the system is actually helping users during the interaction.

This means that future evaluations should expand beyond accuracy rather than replace it. A CRS should be considered successful not only when it recommends relevant options, but also when it helps users understand, explore, and evaluate those options. A more complete evaluation could therefore consider both what the system recommends and how the user experiences the process of reaching that recommendation [6, 9].

4.3 Future Research Directions

The findings of this review raise several questions for future research. First, more research is needed to develop evaluation methods that combine recommendation performance with user experience. Frameworks such as CRS-Que and CONCEPT provide important steps in this direction by considering both system-centered and user-centered factors [6, 9]. However, these approaches also raise the question of how different aspects of user support, such as exploration, understanding, trust, and user control, can be measured consistently across different CRS.

Second, future research should examine CRS over long periods of time. Many CRS studies evaluate users through individual conversations or controlled experiments. However, users' preferences may change over time, and their relationship with a system may also change through repeated interactions. A system that works well during one conversation may not necessarily continue to support the same user effectively over a longer period. Long-term studies could therefore provide a better understanding of how CRS adapt to changing preferences and whether users continue to find these systems useful.

Another important question is how proactive a CRS should be. The literature suggests that introducing unfamiliar options and helping users explore can provide value beyond simply matching an existing preference [10, 11, 12]. However, there needs to be a balance between

helping users discover new possibilities and influencing their decisions too much. This raises the question of whether different users may prefer different levels of system guidance. Research on conversational styles already suggests that users can have different preferences for how they interact with CRS [13]. Future research could therefore examine how the level of proactivity should be adapted to individual users.

Finally, the increasing use of Large Language Models presents both opportunities and challenges. Although LLM-based conversational recommender systems provide more neutral, flexible, and engaging interactions, transparency, factual accuracy, privacy, and explainability remain issues to solve [18, 24]. Therefore, future studies should explore methods for balancing conversational flexibility with trustworthy and transparent recommendations while maintaining positive user experiences.

Overall, the findings suggest that the future of CRS should not be measured only by how accurately a system can recommend an item. A CRS can also support users in discovering, understanding, and evaluating their options. This does not make recommendation accuracy less important. Instead, it places accuracy within a larger goal of helping users make informed decisions while maintaining control over the recommendation process.

5 Acknowledgement

I would like to thank Manveer Kalirai for helping improve this work.

6 Conclusion

Conversational Recommender Systems (CRSs) have evolved from traditional recommender systems that relied on users' historical data and one-shot recommendations. Through conversational interactions, modern CRSs gather user information over time, adapt their recommendations, and provide a more personalized experience. Recent advancements in Large Language Models have further improved the flexibility and naturalness of conversations, enabling the system to understand the users' intentions better and maintain multi-turn dialogues.

This literature review identified three major research directions in recent CRS research. First, researchers have been constantly improving recommendation performance through strategic planning, reinforcement learning, adaptive user modeling, and personalized dialogue strategies. Second, more attention is being paid to supporting human information exploration by improving discoverability, providing richer explanations, and encouraging users to explore unfamiliar options rather than simply reinforcing existing preferences. Finally, research has gradually shifted from focusing solely on recommendation accuracy toward human-centered measures that consider trust, explainability, conversational quality, and overall user experience.



Overall, we suggest that the role of Conversational Recommender Systems is evolving beyond providing accurate recommendations. Rather than acting only as a recommendation tool, CRSs are increasingly designed to serve as intelligent conversational partners that help users discover information, refine their preferences, and make informed decisions. As conversational technologies continue to advance, future research should continue emphasizing human-centered design principles to ensure that conversational recommender systems are not only about accuracy but also trustworthy, transparent, and genuinely supportive of users throughout their conversation with the systems to find information.

References

- [1] Christakopoulou, K., Radlinski, F., & Hofmann, K. (2016). Towards Conversational Recommender Systems. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. <https://dl.acm.org/doi/epdf/10.1145/2939672.2939746>
- [2] Dai, X., Wang, Z., Xie, J., Yu, T., & Lui, J. C. (2024). Online learning and detecting corrupted users for conversational recommendation systems. IEEE Transactions on Knowledge and Data Engineering, 36(12), 8939-8953.
<https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=10643701>
- [3] David Moher, Alessandro Liberati, Jennifer Tetzlaff, and Douglas G. Altman. 2009. Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement. BMJ 339 (July 2009), b2535. <https://www.bmj.com/content/339/bmj.b2535> Publisher: British Medical Journal Publishing Group Section: Research Methods& Reporting
- [4] Du, H., Peng, B., & Ning, X. (2025, April). SAPIENT: Mastering multi-turn conversational recommendation with strategic planning and Monte Carlo tree search. In 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) (pp. 2629-2648).
<https://aclanthology.org/2025.naacl-long.133.pdf>
- [5] Hidayat, M. L., Ahmad, I., Assidiqi, M. H., Khanzada, T. J., & Meccawy, M. (2025, May). Knowledge Discovery and Scientometric Analysis of Conversational Artificial Intelligence (CAI) for Mental Health. In 2025 International Conference on Innovation in Artificial Intelligence and Internet of Things (AIIT) (pp. 1-15). IEEE.
[google.com/url?q=https://www.researchgate.net/publication/393929959_Knowledge_Discovery_and_Scientometric_Analysis_of_Conversational_Artificial_Intelligence_CAI_for_Mental_Health&sa=D&source=docs&ust=1791158695782056&usg=AOvVaw1HOEjoLCPiL9WEyziTZVgW](https://www.researchgate.net/publication/393929959_Knowledge_Discovery_and_Scientometric_Analysis_of_Conversational_Artificial_Intelligence_CAI_for_Mental_Health&sa=D&source=docs&ust=1791158695782056&usg=AOvVaw1HOEjoLCPiL9WEyziTZVgW)
- [6] Huang, C., Qin, P., Deng, Y., Lei, W., Lv, J., & Chua, T. S. (2024). Concept--An Evaluation Protocol on Conversational Recommender Systems with System-centric and User-centric Factors. arXiv preprint arXiv:2404.03304. <https://arxiv.org/pdf/2404.03304>
- [7] Jannach, D., Manzoor, A., Cai, W., & Chen, L. (2020). A Survey on Conversational Recommender Systems. ACM Computing Surveys (CSUR), 54, 1 - 36.
<https://dl.acm.org/doi/epdf/10.1145/3453154>

-
- [8] Jin, Y., Chen, L., Cai, W., & Pu, P. (2021). Key Qualities of Conversational Recommender Systems: From Users' Perspective. Proceedings of the 9th International Conference on Human-Agent Interaction. <https://dl.acm.org/doi/epdf/10.1145/3472307.3484164>
- [9] Jin, Y., Chen, L., Cai, W., & Zhao, X. (2023). CRS-Que: A User-centric Evaluation Framework for Conversational Recommender Systems. ACM Transactions on Recommender Systems, 2, 1 - 34. <https://dl.acm.org/doi/epdf/10.1145/3631534>
- [10] Kalirai, M., & Kuzminykh, A. (2024, July). You Today, Better Tomorrow: Envisioning the Role of Conversation in Recommender Systems of the Future. In Proceedings of the 6th ACM Conference on Conversational User Interfaces (pp. 1-5). <https://dl.acm.org/doi/fullHtml/10.1145/3640794.3665881>
- [11] Kalirai, M., & Kuzminykh, A. (2026). Explanation Driving Exploration: Aligning Conversational Recommender Systems with Users' Exploratory Information Needs. Proceedings of the 31st International Conference on Intelligent User Interfaces. <https://dl.acm.org/doi/10.1145/3742413.3789165>
- [12] Kalirai, M., Williams, A. C., & Kuzminykh, A. (2024, March). Toward Faceted Skill Recommendation in Intelligent Personal Assistants. In Proceedings of the 29th International Conference on Intelligent User Interfaces (pp. 640-649). <https://dl.acm.org/doi/pdf/10.1145/3640543.3645201>
- [13] Kostic, I., Balog, K., & Gadiraju, U. (2025). Should We Tailor the Talk? Understanding the Impact of Conversational Styles on Preference Elicitation in Conversational Recommender Systems. Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization. <https://dl.acm.org/doi/epdf/10.1145/3699682.3728353>
- [14] Lei, W., He, X., de Rijke, M., & Chua, T. (2020). Conversational Recommendation: Formulation, Methods, and Evaluation. Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval. <https://hexiangnan.github.io/papers/sigir20-tutorial.pdf>
- [15] Mahmud, R., Berkovsky, S., Prasad, M., & Kocaballi, A.B. (2025). Understanding User Preferences for Interaction Styles in Conversational Recommender System: The Predictive Role of System Qualities, User Experience, and Traits. Proceedings of the 37th Australian Conference on Human-Computer Interaction. <https://dl.acm.org/doi/epdf/10.1145/3764687.3764722>

- [16] Manzoor, A., Cai, W., & Jannach, D. (2023). Factors Influencing the Perceived Meaningfulness of System Responses in Conversational Recommendation. IntRS@ RecSys, 23. https://web-ainf.aau.at/pub/jannach/files/Workshop_IntRS_2023.pdf
- [17] McConvey, K., Guha, S., & Kuzminykh, A. (2023). A Human-Centered Review of Algorithms in Decision-Making in Higher Education. Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. <https://dl.acm.org/doi/epdf/10.1145/3544548.3580658>
- [18] Qin, P., Huang, C., Deng, Y., Lei, W., & Chua, T. (2024). Beyond Persuasion: Towards Conversational Recommender System with Credible Explanations. ArXiv, abs/2409.14399. https://ink.library.smu.edu.sg/cgi/viewcontent.cgi?article=10616&context=sis_research
- [19] Shen, Q., Wu, L., Zhang, Y., Pang, Y., Wei, Z., Xu, F., ... & Pei, J. (2024). Multi-interest multi-round conversational recommendation system with fuzzy feedback based user simulator. ACM Transactions on Recommender Systems, 2(4), 1-29. <https://dl.acm.org/doi/pdf/10.1145/3616379>
- [20] Shen, X., Lee, D., Ranjan, S., Harsha, S.S., Sevak, P., & Li, Y. (2024). Enhancing Discoverability in Enterprise Conversational Systems with Proactive Question Suggestions. ArXiv, abs/2412.10933. <https://arxiv.org/pdf/2412.10933>
- [21] Wei, C.Z., Kim, Y., & Kuzminykh, A. (2023). The Bot on Speaking Terms: The Effects of Conversation Architecture on Perceptions of Conversational Agents. Proceedings of the 5th International Conference on Conversational User Interfaces. <https://dl.acm.org/doi/epdf/10.1145/3571884.3597139>
- [22] Wen, B., Bu, X., & Shah, C. (2022). EGCR: Explanation Generation for Conversational Recommendation. ArXiv, abs/2208.08035. <https://arxiv.org/html/2208.08035v2>
- [23] Xia, Y., Wu, J., Yu, T., Kim, S., Rossi, R. A., & Li, S. (2023, August). User-regulation deconfounded conversational recommender system with bandit feedback. Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining (pp. 2694-2704). <https://dl.acm.org/doi/pdf/10.1145/3580305.3599539>
- [24] Yun, S., & Lim, Y. (2025). User Experience with LLM-powered Conversational Recommendation Systems: A Case of Music Recommendation. Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems. <https://dl.acm.org/doi/epdf/10.1145/3706598.3713347>



[25] Zhao, X., Yan, M., Zhang, Y., Deng, Y., Wang, J., Zhu, F., ... & Chua, T. S. (2025). Reinforced Strategy Optimization for Conversational Recommender Systems via Network-of-Experts. arXiv preprint arXiv:2509.26093. <https://arxiv.org/pdf/2509.26093>