

Symbiotic Conversational Recommender Systems: A New Approach to Improving Transparency and Persuasion

Alessandro Petruzzelli

¹ Università degli studi di Bari "Aldo Moro"

² `alessandro.petruzzelli@uniba.it`

Abstract. This project proposal explores the concept of human-machine symbiosis within conversational recommender systems (CRSs), aiming to develop CRSs that are more transparent, and persuasive and can better support users in decision-making tasks. One way to equip CRSs with these capabilities is by making use of the potential of large language models (LLMs). The research focuses on three main research directions: (1) integrating knowledge into LLMs to optimize their recommendation capabilities in CRSs; (2) exploring new ways to fine-tune a pre-trained model for conversational recommendations, without forgetting the knowledge it has already learned; (3) evaluating the impact of an LLM-based CRS on users in terms of transparency, engagement, and persuasion. Through this research, the aim is to overcome limitations in current CRS approaches and develop a more collaborative and user-centric recommendation system.

1 Introduction

In recent years, the concept of human-machine symbiosis has gained prominence, as AI systems have become increasingly integrated into our daily lives. This shift has led to a growing interest in creating collaborative environments where humans and AI can work together. One crucial aspect of this collaboration is the ability of AI systems to explain their actions to human collaborators, bridging the gap between the AI's model and the human's mental model [12]. One of the processes that could benefit from symbiotic collaboration is decision-making. The machine can help the human to make better decisions by providing them with information and suggesting solutions. The human can help the machine to make better decisions by providing it with context, explaining the rationale behind their decisions, and correcting any errors. Recommender Systems can be considered as a form of collaboration between humans and AI, which aim to provide recommendations and help users during the decision-making process [3].

In this research field, some proposed solutions move towards human collaboration by interacting with them using natural language. These solutions, called conversational recommender systems (CRSs), are able to understand the user's needs and goals and provide recommendations by interacting through multi-turn dialog [11]. To bridge the gap between the AI model and the human mental

model, and to increase transparency and persuasion, CRSs need to be able to act in a more natural way. They should allow users to express their preferences freely and provide and explain recommendations effectively. This is where the concept of symbiotic AI comes in. Symbiotic AI is a type of AI that is designed to work in collaboration with humans. CRSs that are based on symbiotic AI can explain their recommendations to users in a way that is understandable and transparent.

This project proposal explores the limitations of the current state of the art of conversational recommender systems (CRSs) and proposes a novel approach that combines symbiotic AI with recommender systems. The goal of this Ph.D. project is to investigate the role of large language models (LLMs) in improving the user experience of CRSs by focusing on three main research questions: (1) *How can knowledge be instilled in LLMs to optimize their performance in CRSs?* (2) *How do LLMs perform in the context of conversational recommendation?* (3) *How do LLM-based CRSs compare to traditional recommender systems in terms of user engagement, transparency and the persuasion of the users?*

2 Background about Symbiotic AI

Human-machine symbiosis is a collaborative partnership between humans and AI systems that can correct individual shortcomings and collectively overcome limitations through mutual enhancement. As AI advances, there is a growing emphasis on the comprehensibility and clarity of AI-assisted tasks. A challenge in designing more interpretable and reliable AI systems is incorporating the human mental model. This is because the interpretability of behavior is dependent on its alignment with human expectations [2].

On the flip side, to achieve a truly symbiotic collaboration between humans and AI, users need to be more informed about the AI model’s capabilities and provide it with relevant information and context. This collaborative process, which is now a common practice in the development of most AI models, can be made more natural through an iterative cycle where both humans and AI systems actively contribute to reaching their shared goals.

An innovative symbiotic approach to recommender systems has been suggested in [1]. This conceptual framework differs from the traditional human-centered design approach in favor of a symbiotic paradigm, in which humans and the system exchange knowledge in a way that benefits both of them.

3 Related Work

This section discusses the current state-of-the-art Conversational Recommender Systems (CRS) and examines their limitations in terms of user-friendliness and interactive capabilities. This exploration highlights the need for a new perspective, which led to the introduction of an innovative symbiotic approach. Two main approaches of CRSs are presented: (1) solutions based on statically defined domain knowledge, and (2) end-to-end learning models.

3.1 Statically Defined Domain Knowledge CRS

An essential characteristic of a Statically Defined Domain Knowledge CRS lies in its architecture, which consists of three components: (1) The Dialog Manager, which controls the entire interaction process, including recognizing user intent and expressed preferences and generating appropriate responses. (2) The User Modeling System, which is responsible for formulating user preferences in a way that is consistent with the model’s representation scheme. (3) The Recommender Engine, which serves as the recommendation model and suggests items. These three constituents establish connections with a Background Knowledge Base, which is a fixed component that encapsulates and models domain information. This conceptual framework differs from the traditional human-centered design approach in favor of a symbiotic paradigm, in which humans and the system exchange knowledge in a way that benefits both of them.

This fixed architecture restricts how users can express their preferences. Some solutions adopt the System Ask User Respond (SAUR) paradigm [16], in which users can only respond to questions initiated by the system. These questions are typically formulated using predefined language patterns or templates and are usually about specific things in the domain.

To address this limitation and improve the user experience, some solutions have been developed that combine the features of digital assistants with CRS [9]. Nevertheless, the architecture still limits users to expressing preferences about predetermined entities, preventing nuanced preference elicitation. In the movies domain, one approach integrates review-derived aspects into domain knowledge and movie representations [14], yet it does not fully address this issue. In fact, all the existing solutions are applied on the user side. This means that the recommender responses are based on fixed templates. This can lead to a poor user experience, which could be improved with more accurate signals from the model [15].

3.2 End-To-End Learning CRS

The research field of CRSs has moved to an end-to-end approach since the introduction of the dialog dataset Redial [13]. The promise of this new approach is to overcome all the components of statically defined domain knowledge CRSs by using a deep learning model to analyze user input, model their preferences, provide a recommendation, and generate a natural language response.

However, this new approach has some shortcomings. Firstly, most solutions [4, 17] employ a preprocessing of the sequences to detect user-expressed preferences, which inherits the limitations of statically defined domain knowledge CRSs. This shift in focus towards the recommendation process overlooks the crucial role of the user during the interaction. These solutions model user preferences based on a single iteration, but often recommendations are refined through an iterative feedback loop where the model assists the user in expressing their preferences, providing examples and explanations.

The generative abilities of these models have been tested in a human evaluation conducted by Jannach and Manzoor [10]. They found that the generated responses were of poor quality, leading to a bad user experience and usability. This was because more than two-thirds of the responses were identical to those in the training set. The authors suggest that the evaluation methodology should be changed from metrics that measure user satisfaction and usability, as these metrics can be misleading. For example, they proved that a model with a higher accuracy metric may actually have worse performance than a model with a lower accuracy metric.

4 Research Objectives

This section proposes a novel approach to CRSs that leverages the foundations of symbiotic AI. In the context of CRSs, this means that the AI system and the user would work together to find the best recommendations for the user. To achieve symbiosis, the focus is placed on conversation and explanation. The AI system would engage in a conversation with the user to understand their preferences and needs. The AI system would also explain its recommendations to the user so that the user can understand why the AI system made those recommendations.

Large language models (LLMs) can be used to implement this approach. LLMs can extract high-quality representations of textual features and leverage extensive external knowledge [5]. This can help to overcome the limitations of statically defined domain knowledge CRS. Specifically, LLMs can overcome the fixed representation of domain knowledge by encoding it in a more flexible and expressive way. This allows users to express their preferences more freely, which helps the model to better understand their mental model.

Recent work has proposed LLM-based solutions for CRS [7, 8]. These solutions can be considered as end-to-end, as they use the generative capabilities of LLMs to provide accurate recommendations without exploring their conversational abilities.

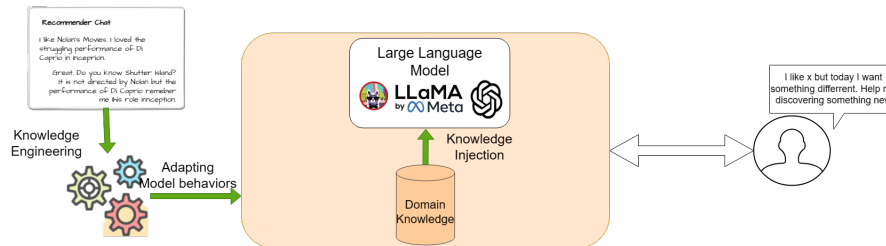


Fig. 1. High-level architecture of a symbiotic CRS

The objective of this PhD project is to propose a conceptual framework for a symbiotic CRS, whose high-level architecture is shown in Fig. 1. Specifically, the research will focus on addressing the following research questions:

(RQ1) How can knowledge be instilled in LLMs to optimize their performance in CRSs? One way to instill knowledge in LLMs to optimize their performance in CRSs is through fine-tuning methods. Fine-tuning is the process of adjusting the parameters of a pre-trained model to improve its performance on a specific task. This can be done by training the model on a smaller, task-specific dataset, allowing it to learn and adapt to the nuances of the new task. This step in the recommendation process involves identifying the domain knowledge that should be used to generate recommendations. This knowledge is essential for providing explanations, as the model must be able to highlight the similarities between the user’s preferences and the recommended items.

(RQ2) How do LLMs perform in the context of conversational recommendation? Large language models (LLMs) are often fine-tuned to adapt to new contexts. However, fine-tuning an LLM multiple times can lead to the catastrophic forgetting problem, where the model loses its ability to perform well on previous tasks [6]. To address this challenge, no effective solutions have been proposed. The aim of this project is to explore solutions to overcome this issue and allow the model to act like a conversational recommender system (CRS) by exploiting its memorized knowledge.

(RQ3) How do LLM-based CRSs compare to traditional recommender systems in terms of transparency, user engagement the persuasion of the users? This RQ follows what is proposed in [10] which highlights the importance of user-centric metrics. Transparency can be achieved through explainable AI (XAI), which allows users to understand why certain recommendations are made. This can be done by providing explanations in natural language or by visualizing the decision-making process. User engagement can be enhanced by the chat ability of the model. Users are more likely to continue a conversation if the model provides interesting suggestions. The concept of "interesting" can be broadened beyond accuracy to include serendipity and novelty. The last aspect of CRSs that has not been considered in the proposed solution is persuasion.

5 Conclusion

The research objectives outlined in this project proposal address the challenges and opportunities in the context of CRS. The objectives are to instill knowledge in LLMs, optimize conversation recommendation ability, and evaluation user-center metrics. By addressing these objectives, I move closer to the realization of a symbiotic relationship between humans and AI in recommendation systems, which promises improved user experiences and outcomes.

Acknowledgment

I would like to express my sincere gratitude to the SWAP³ group, led by Professor Semeraro, for their support and guidance in the development of this project proposal. I am particularly grateful to Professor Musto and Professor de Gemmis for their insights and expertise.

References

1. Brusilovsky, P., de Gemmis, M., Felfernig, A., Lops, P., Polignano, M., Semeraro, G., Willemsen, M.C.: Joint workshop on interfaces and human decision making for recommender systems (intra'22). p. 667–670. <https://doi.org/10.1145/3523227.3547413>
2. Chakraborti, T., Kulkarni, A., Sreedharan, S., Smith, D.E., Kambhampati, S.: Explicability? legibility? predictability? transparency? privacy? security? the emerging landscape of interpretable agent behavior. In: ICAPS '19
3. Chen, L., de Gemmis, M., Felfernig, A., Lops, P., Ricci, F., Semeraro, G.: Human decision making and recommender systems. *ACM Trans. Interact. Intell. Syst.* (2013), <https://doi.org/10.1145/2533670.2533675>
4. Chen, Q., Lin, J., Zhang, Y., Ding, M., Cen, Y., Yang, H., Tang, J.: Towards knowledge-based recommender dialog system. In: EMNLP-IJCNLP '19'. pp. 1803–1813. <https://aclanthology.org/D19-1189>
5. Chen, Z., Mao, H., Li, H., Jin, W., Wen, H., Wei, X., Wang, S., Yin, D., Fan, W., Liu, H., Tang, J.: Exploring the potential of large language models (llms) in learning on graphs (2023)
6. He, T., Liu, J., Cho, K., Ott, M., Liu, B., Glass, J., Peng, F.: Analyzing the forgetting problem in the pretrain-finetuning of dialogue response models (2021)
7. Hou, Y., Zhang, J., Lin, Z., Lu, H., Xie, R., McAuley, J., Zhao, W.X.: Large language models are zero-shot rankers for recommender systems (2023)
8. Huang, X., Lian, J., Lei, Y., Yao, J., Lian, D., Xie, X.: Recommender ai agent: Integrating large language models for interactive recommendations (2023)
9. Iovine, A., Narducci, F., Semeraro, G.: Conversational recommender systems and natural language: A study through the converse framework. *Decision Support Systems* (2020). <https://doi.org/https://doi.org/10.1016/j.dss.2020.113250>
10. Jannach, D., Manzoor, A.: End-to-end learning for conversational recommendation: A long way to go? In: IntRS@RecSys '20'
11. Jannach, D., Manzoor, A., Cai, W., Chen, L.: A survey on conversational recommender systems. *ACM Comput. Surv.* **54**(5) (2021). <https://doi.org/10.1145/3453154>
12. Jarrahi, M.H.: Artificial intelligence and the future of work: Human-ai symbiosis in organizational decision making. *Business Horizons* **61**(4), 577–586 (2018). <https://doi.org/https://doi.org/10.1016/j.bushor.2018.03.007>
13. Li, R., Kahou, S., Schulz, H., Michalski, V., Charlin, L., Pal, C.: Towards deep conversational recommendations. In: NIPS '18. p. 9748–9758
14. Martina, A.F.M., Musto, C., Iovine, A., de Gemmis, M., Narducci, F., Semeraro, G.: A virtual assistant for the movie domain exploiting natural language preference elicitation strategies. In: UMAP '22. p. 7–12. <https://doi.org/10.1145/3511047.3536407>

³ <http://www.di.uniba.it/~swap/>

15. Radensky, M., Séguin, J.A., Lim, J.S., Olson, K., Geiger, R.: “i think you might like this”: Exploring effects of confidence signal patterns on trust in and reliance on conversational recommender systems. In: FAccT '23. p. 792–804. <https://doi.org/10.1145/3593013.3594043>
16. Zhang, Y., Chen, X., Ai, Q., Yang, L., Croft, W.B.: Towards conversational search and recommendation: System ask, user respond. In: CIKM '18. pp. 177–186
17. Zhou, K., Zhao, W.X., Bian, S., Zhou, Y., Wen, J.R., Yu, J.: Improving conversational recommender systems via knowledge graph based semantic fusion. <https://doi.org/10.1145/3394486.3403143>