Enhancing Healthcare Decisions with Explainable Session-Based Recommendations

Jiahao Zhu*

Beijing Haidian District Private New Oriental School, Beijing, China

*zhu17310852891@163.com

Abstract. Recommender systems have demonstrated efficacy in various domains, but healthcare poses unique challenges with its complex medical knowledge and demand for explainable recommendations. This paper introduces a novel session-based, reasoning-focused recommender system for healthcare professionals during patient consultations. The system utilizes a hybrid knowledge graph integrating medical ontologies, real-time patient data, and external resources. Advanced reasoning over this graph infers diagnoses, predicts drug interactions, and suggests treatments, particularly effective in cold-start scenarios. It dynamically updates suggestions during consultations as new data emerges, prioritizing transparency by providing clear explanations for recommendations. Evaluation on medical datasets and user studies shows improved diagnostic accuracy, faster diagnosis times, and high satisfaction with explainable recommendations. This research advances knowledge-based reasoning in healthcare, enhancing decision-making and patient care.

Keywords: Recommender Systems, Knowledge Graphs, Healthcare, Explainable AI, Session-Based Recommendation, Cold-Start Problem.

1 Introduction

Recommender systems have become integral to modern information consumption, guiding users towards items (products, movies, content) they might enjoy based on their past preferences and behavior [8]. These systems have proven particularly effective in domains like e-commerce and entertainment, but their application in healthcare remains a challenging frontier.

The healthcare domain presents unique challenges for recommender systems due to the complexity of medical knowledge, the necessity for explainable and contextually relevant recommendations, and the prevalence of cold-start scenarios. Traditional systems struggle with the intricate nature of medical information and must ensure recommendations are accurate, transparent, and understandable for patient safety. To address these challenges, we propose a novel system leveraging a reasoning-enabled knowledge graph for explainable, session-based recommendations during consultations. This ap-

L. Shi et al. (eds.), *Proceedings of the 4th International Conference on Public Administration, Health and Humanity Development (PAHHD 2024)*, Advances in Social Science, Education and Humanities Research 868,
https://doi.org/10.2991/978-2-38476-295-8_9

proach incorporates rich medical knowledge and advanced reasoning mechanisms, enhancing decision-making and patient care, particularly in cold-start situations, aligning with recent trends in knowledge graph integration [3].

Despite progress in recommender systems, challenges persist in healthcare due to the unique nature of medical knowledge and decision-making. Adoption of knowledge graphs (KGs) is limited by the complexity of maintaining comprehensive medical KGs [3]. Existing systems often lack transparency, hindering trust and adoption [8]. Many systems fail to handle dynamic healthcare consultations, resulting in static recommendations. The cold-start problem, where new patients or rare conditions lack historical data, is significant. Evaluating these systems is challenging due to the sensitivity of medical data. Addressing these issues requires systems leveraging medical KGs, providing transparent explanations, adapting to dynamic contexts, and handling cold-start scenarios.

To overcome challenges in healthcare recommendation, we propose a novel framework combining knowledge graph-based reasoning and session-based recommendation. Our system constructs a comprehensive hybrid knowledge graph (KG) integrating diverse medical sources, such as medical ontologies (e.g., SNOMED CT [8]), real-time patient data, and external resources. Advanced reasoning techniques infer diagnoses, predict drug interactions, and suggest treatments, uncovering complex relationships. The session-based approach continuously updates recommendations during patient consultations, ensuring relevant suggestions. Clear explanations detailing the reasoning process enhance transparency and trust. Inspired by prior work [7], our framework aims to improve clinical recommendations and patient care.

2 Related Work

Recommender systems play a pivotal role in healthcare by aiding clinical decision-making through personalized recommendations [14]. These systems utilize various approaches such as content-based filtering, collaborative filtering, and hybrid methods to enhance accuracy and recommendation diversity [14]. Knowledge graph (KG)-based recommendations are particularly effective, integrating structured medical knowledge to improve explainability and address challenges like cold-start scenarios [3]. They leverage KG embeddings (KGE) such as TransE, DistMult, and ComplEx to represent entities and relationships, enhancing recommendation models like matrix factorization and neural networks [10]. Path-based methods and graph neural networks (GNNs) further refine KG-based recommendations by capturing complex relationships and improving reasoning capabilities [6]. Our proposed system integrates KG-based reasoning with session-based recommendation (SBR) frameworks to address dynamic contexts and enhance decision support.

KGs provide a structured framework to encode relationships between medical entities, crucial for accurate and contextually relevant recommendations in healthcare [12] [13]. Path-based reasoning extracts implicit knowledge from KGs, aiding in diagnosis prediction and treatment recommendations by uncovering intricate connections from patient data [13]. Embedding-based approaches predict missing links within KGs, while neural logic reasoning combines symbolic logic with neural networks for logical inference, essential in managing uncertainty and incomplete data in medical contexts. Reinforcement learning (RL) techniques further optimize KG-based reasoning by training agents to navigate complex KG structures, enhancing recommendation accuracy [13].

Session-based recommendation (SBR) complements KG-based approaches by predicting short-term user preferences during consultations, adapting recommendations in real-time as new data emerges [9]. In healthcare, SBR frameworks like recurrent neural networks (RNNs), attention mechanisms, and graph neural networks (GNNs) are pivotal for adapting to rapidly changing patient needs and preferences, despite challenges in handling cold-start scenarios. Our research aims to pioneer a healthcare-specific SBR system that integrates KGs to leverage medical knowledge and reasoning capabilities, addressing critical gaps in personalized decision support for healthcare professionals.

3 Methodology

Our reasoning-enabled recommender system uses a robust hybrid knowledge graph (KG) construction approach for effectiveness. We integrate diverse medical knowledge sources to build a comprehensive representation. Firstly, established medical ontologies and knowledge bases like SNOMED CT [8], ICD-10 [14] [15], and RxNorm provide standardized vocabularies and relationships among diseases, symptoms, drugs, and treatments. Secondly, real-time patient data, including medical history, medications, allergies, demographics, and lab or imaging results, personalize recommendations securely and compliantly. Thirdly, external resources like clinical trial outcomes, drug databases, and medical guidelines ensure the KG is current with the latest medical insights. Our strategy includes entity linking techniques [2] and knowledge graph completion methods to align entities, resolve inconsistencies, and maintain data quality. This hybrid KG empowers the reasoning engine to derive insights and generate personalized, contextually aware recommendations effectively.

Our system's core is the reasoning engine, deriving insights from a hybrid knowledge graph (KG) using knowledge graph embedding, graph neural networks (GNNs), and rule-based reasoning. GNNs propagate information for diagnosis inference, capturing relationships among symptoms, diseases, and risk factors to suggest diagnoses based on patient-specific data [4] [5] [11]. Rule-based reasoning ensures diagnoses align with medical guidelines. Predicting drug interactions involves knowledge graph embeddings modeling drugs and interactions, augmented by domain-specific rules and external databases [1]. For treatment pathways, the engine uses knowledge-aware graph traversal and reinforcement learning (RL) to propose personalized plans based on patient conditions, history, and preferences, ensuring contextually appropriate and evidence-based recommendations.

Our session-based recommendation framework provides dynamic, contextually relevant guidance to healthcare professionals during patient consultations. It begins by initializing a session with pertinent patient data from the electronic health record (EHR), integrating medical history, medications, allergies, demographics, and recent test results into the knowledge graph (KG). Concurrently, relevant medical knowledge is retrieved based on the patient's concerns. As the consultation progresses, clinicians input new data, such as symptoms or test outcomes, updated in real-time within the session's KG. The reasoning engine uses inference techniques to generate ongoing recommendations—diagnoses, tests, or treatments—continuously refining suggestions as the patient's condition evolves. Reinforcement learning (RL) optimizes recommendation strategies based on feedback, adapting to each consultation phase.

4 Experimental Analysis

4.1 Experimental Setup

To evaluate the effectiveness of our session-based reasoning recommender system in healthcare, we utilize two distinct datasets: a curated patient case dataset sourced from electronic health records (EHRs), de-identified for privacy protection, and simulated patient scenarios. The curated dataset includes essential patient information such as diagnoses, symptoms, medications, and lab results, serving as the benchmark for evaluating diagnostic accuracy and treatment recommendations. These scenarios allow us to assess the system's adaptability and real-time recommendation capabilities under controlled conditions. Our evaluation metrics encompass accuracy metrics like Accuracy@K, Precision@K, Recall@K, F1 Score@K, and Mean Reciprocal Rank, which measure ranking effectiveness. Explainability is evaluated through metrics such as Explanation Precision@K and Explanation Recall@K, complemented by qualitative assessments from healthcare professionals regarding clarity and usefulness of explanations. We benchmark our system against rule-based methods and traditional recommendation models like Collaborative Filtering and Matrix Factorization to demonstrate superior performance in accuracy, explainability, and recommendation novelty and diversity. This comprehensive approach ensures robust evaluation across diverse clinical contexts, enhancing decision support in healthcare.

4.2 Offline Evaluation

Table 1. Diagnosis Prediction Accuracy

Model	Accuracy@1	Accuracy @3	Accuracy @5	MRR
Proposed System	0.85	0.92	0.95	0.88
Rule-Based System	0.78	0.87	0.91	0.81
Collaborative Filtering	0.72	0.83	0.88	0.75
Matrix Factorization	0.70	0.81	0.86	0.73

Table 2. Treatment Recommendation Accuracy

Model	Accuracy @1	Accuracy @3	Accuracy @5	MRR
Proposed System	0.82	0.89	0.93	0.85
Rule-Based System	0.75	0.84	0.89	0.78
Collaborative Filtering	0.68	0.79	0.85	0.71
Matrix Factorization	0.65	0.76	0.82	0.68

We conducted offline experiments on the curated patient case dataset to evaluate the accuracy, novelty, and diversity of our proposed system's recommendations, comparing its performance against the rule-based system and traditional recommendation models (Collaborative Filtering and Matrix Factorization). For diagnosis prediction accuracy, Table 1 shows our system consistently outperforms the baselines across all metrics, achieving Accuracy@1 of 0.85, Accuracy@3 of 0.92, Accuracy@5 of 0.95, and an MRR of 0.88. Similarly, for treatment recommendation accuracy, Table 2 illustrates our system's superior performance with Accuracy@1 of 0.82, Accuracy@3 of 0.89, Accuracy@5 of 0.93, and an MRR of 0.85 compared to the baselines. Furthermore, we evaluated the novelty and diversity of recommendations, where our system recommends less popular items (average popularity of 0.3) and more diverse items (average dissimilarity of 0.7) compared to traditional models. These results demonstrate that our system effectively leverages medical knowledge and reasoning to provide accurate, novel, and diverse recommendations tailored to patient-specific contexts in healthcare settings.

4.3 User Study

After each consultation, participants were asked to rate their satisfaction with the system and the quality of its explanations. The results are summarized in Table 3.

Table 3. User Satisfaction and Explanation Quality

Metric	Percentage
Overall Satisfaction (System is helpful)	85%
Clarity of Explanations	78%
Comprehensiveness of Explanations	82%
Usefulness of Explanations	80%

To assess the real-world impact of our system, we conducted a user study with healthcare professionals who interacted with patient scenarios using our system to aid in diagnosis and treatment recommendations. Participants' performance and satisfaction were compared with and without the system's assistance. We measured the time taken by participants to reach a diagnosis and recommend treatment, illustrating a 20% reduction in diagnosis time and a 15% decrease in treatment recommendation time when using the system. Moreover, diagnostic accuracy significantly improved with the system, showing an average increase of 10% across scenarios. After each consultation, participants rated their satisfaction and the quality of system explanations, with 85% reporting overall satisfaction and high ratings for clarity (78%), comprehensiveness (82%), and usefulness (80%) of explanations, as summarized in Table 3.

5 Conclusion

Through this study, we have developed a session-based reasoning recommendation system that significantly improves the accuracy and efficiency of medical recommendations. The system utilises hybrid knowledge graphs and advanced reasoning techniques to provide contextually relevant and interpretable recommendations, providing strong support for healthcare professionals in complex clinical environments. Experimental results show that the system excels in several aspects, including diagnostic accuracy, time efficiency and user satisfaction.

References

1. Y. Cao, L. Hou, J. Li, and Z. Liu. Neural collective entity linking. In *Proceedings of COLING 2018*, pages 675–686, 2018.
2. Q. Guo, F. Zhuang, C. Qin, H. Zhu, X. Xie, H. Xiong, and Q. He. A survey on knowledge graph-based recommender systems. *IEEE Transactions on Knowledge and Data Engineering*, 34(8):3549–3568, 2020.
3. W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *Neural Information Processing Systems*, volume 30, 2017.
4. B. Hui, D. Yan, H. Chen, and W.-S. Ku. Time-sensitive poi recommendation by tensor completion with side information. In *2022 IEEE 38th International Conference on Data Engineering*, pages 205–217. IEEE, 2022.
5. F. O. Isinkaye, Y. O. Folajimi, and B. A. Ojokoh. Recommendation systems: Principles, methods and evaluation. *Egyptian informatics journal*, 16(3):261–273, 2015.
6. S. Ji, S. Pan, E. Cambria, P. Marttinen, and S. Y. Philip. A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2):494–514, 2021.
7. T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2016.
8. W. Krichene and S. Rendle. On sampled metrics for item recommendation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1748–1757, 2020.
9. Y. Zhang and Q. Yao. Knowledge graph reasoning with relational digraph. In *Proceedings of The Web Conference 2022*, pages 912–924, 2022.
10. Y. Zhang, Q. Yao, L. Yue, X. Wu, Z. Zhang, Z. Lin, and Y. Zheng. Emerging drug interaction prediction enabled by flow-based graph neural network with biomedical network. *Nature Computational Science*, 2023.
11. X. Wang, T. Jia, C. Wang, K. Xu, Z. Shu, K. Yang, and X. Zhou. Knowledge graph completion based on tensor decomposition for disease gene prediction. *arXiv preprint arXiv:2302.09335*, 2023.

12. S. Wang, H. Li, C. C. Cao, X.-H. Li, and N. N. Fai. Tower bridge net (tb-net): Bidirectional knowledge graph aware embedding propagation for explainable recommender systems. In *2022 IEEE 38th International Conference on Data Engineering*, pages 3268–3279. IEEE, 2022.
13. B. Hui, D. Yan, H. Chen, and W.-S. Ku. Time-sensitive poi recommendation by tensor completion with side information. In *2022 IEEE 38th International Conference on Data Engineering*, pages 205–217. IEEE, 2022.
14. S. Ji, S. Pan, E. Cambria, P. Marttinen, and S. Y. Philip. A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2):494–514, 2021.
15. X. Wang, T. Huang, D. Wang, Y. Yuan, Z. Liu, X. He, and T.-S. Chua. Learning intents behind interactions with knowledge graph for recommendation. In *Proceedings of The Web Conference 2021*, pages 878–887, 2021.