Effective and Transparent Course Recommendation through Causal Reasoning with Language Models

Discussion Paper

Neda Afreen, Ludovico Boratto, Gianni Fenu, Mirko Marras and Alessandro Soccol

Department of Mathematics and Computer Science, University of Cagliari, V. Ospedale 72, 09124 Cagliari, Italy

Abstract

Recommender systems are crucial to support learners through the abundance of available online educational resources. Recent advances in educational recommendation have employed knowledge graphs to enhance both the effectiveness and transparency of recommendations. However, these systems primarily rely on correlational reasoning. Such approaches generate user-aligned suggestions through path-based explanations but often fail to capture the underlying, true causal relationships that drive educational progress and decision-making, which inherently depend on the instantiated knowledge graphs. In this paper, we discuss our ongoing efforts in developing reasoning methods in course recommendation and how augmenting models with causal relationships can transform the way recommendations are generated and explained. We discuss the importance of causal inference for developing effective and transparent systems that can recommend not just what other learners with similar profiles choose, but what a learner should study next based on other covariates such as the learning history and context.

Keywords

Casual Reasoning, Educational Recommender System, Knowledge Graph, Personalization.

1. Introduction

Context. Digital learning platforms have expanded opportunities for learners to achieve academic and personal goals through flexible, self-paced learning. As the number of online resources continues to grow, it imposes cognitive load on learners to choose the most suitable course for them [1, 2, 3]. Recommender systems (RSs) have been developed to address this problem of information overload by helping learners navigate content more effectively [4]. With digital platforms continuously evolving, these educational recommender systems must not only be effective but also transparent, and aligned with learners' educational context [5].

Current Work. Recent educational recommender systems have adopted knowledge graphs (KGs) for modeling relationships among relevant entities, such as learners, courses, concepts, and instructors, thereby enabling explainable recommendations through path-based reasoning methods [6]. Our ongoing studies have shown that empowering language models with KGs offers strong performance across utility, beyond utility, and explainability metrics, even in sparse data scenarios [7]. Specifically, generative models have demonstrated the ability to produce diverse explanations by generating paths over KGs.

Research Gap. Despite achieving strong results in utility, beyond-utility, and explainability metrics, current methods fall short in capturing the causal insights, such as the effect of completing a prerequisite course on mastering a subsequent concept. Existing methods lack the ability to distinguish *why* this recommendation works from *what* actually correlates with positive outcomes. For example, a learner who completes an introductory statistics course and later enrolls in a deep learning course may show a common behavioral pattern, but this does not imply that the first course causally contributes to success in the later. Recommending a course is not simply about interest or similarity; rather, it is a

IIR2025: 15th Italian Information Retrieval Workshop, 3th - 5th September 2025, Cagliari, Italy
neda.afreen@unica.it (N. Afreen); ludovico.boratto@acm.org (L. Boratto); fenu@unica.it (G. Fenu); mirko.marras@acm.org (M. Marras); alessandro.soccol@unica.it (A. Soccol)
0000-0002-4507-3516 (N. Afreen); 0000000260533015 (L. Boratto); 0000000346682476 (G. Fenu); 0000000319896057 (M. Marras); 0009-0006-3148-9760 (A. Soccol)

strategic decision that can influence learners' mastery of key skills and progress in future content [8]. Educational objectives are inherently structured and sequential. Learners advance through stages of knowledge acquisition, where concepts are interdependent and pedagogical constraints shape learning paths [9].

Our Perspective. In this paper, we present our ongoing efforts to move beyond correlation-based reasoning by integrating causal inference into course recommendation. Building on our previous work [7], we argue that truly effective and transparent recommendations should reflect not only patterns of learner behavior but also the underlying cause-effect relationships that drive educational progress. While knowledge-graph-based models have demonstrated strong performance across metrics for recommendation utility and path-based explainability, they primarily rely on correlational signals. We explore how augmenting these models with causal reasoning can better align recommendations with learners' actual needs and readiness. This shift raises a key question: how can causal relationships be reliably integrated and validated in real-world educational settings where ground truth is limited? With this in mind, our contribution in this paper is twofold: (i) we contextualize the importance of causal reasoning in educational recommendation, and (ii) we propose strategies for incorporating causal constraints into knowledge-graph-based systems. We invite the community to contribute concrete ideas, empirical insights, and technical solutions related to modeling assumptions, validation strategies, and design trade-offs for integrating causal reasoning into educational recommendation.

2. Characterization of Causal Reasoning in Education

Educational Significance. Causal reasoning involves identifying, modeling, and applying cause-effect relationships to explain outcomes. Educational processes are inherently causal, where understanding one concept leads to the success of another concept [10]. Skills build progressively, and prior learning influences future effectiveness. Learners progress by mastering foundational concepts before advancing, making the learning path a matter of sequencing [11]. Integrating causal reasoning can enable personalized and actionable explanations. Instead of giving rationales "others took this course", causal explanations clarify why a course will improve learning outcomes. Causal models can delay recommendations if key prerequisites are not mastered, offering alternative paths to cover gaps. Despite its potential, causal reasoning remains underexplored in educational recommender systems. This gap presents an opportunity to rethink how recommendations are generated, and moving toward systems that support not only what is popular or similar, but what really help learners progress.

Beyond Correlation. Existing recommendation systems based on correlation often assume that if "learner X who took Course A also took Course B, then B should be recommended to learner Y currently enrolled in A". While this path may capture popularity, it overlooks whether taking Course A actually causes better performance in Course B. As a result, recommendations based solely on these associations may be misleading, since learning outcomes depend on mastery of foundational concepts and the specific characteristics of each learner [12]. In contrast, causal reasoning evaluates whether a learner is prepared for an advanced Course B, or whether a specific sequence of prior courses will lead to higher learning outcomes. Educational effectiveness requires systems to account for prerequisite knowledge, concept dependencies, and effective learning path. By adding causal reasoning, the recommendations can offer better learning outcomes, build learner trust, and align recommendations with progressive learning.

3. Instilling Causal Inference in Educational Recommendation

Current Recommendation Paradigms. In our prior work, we explored how knowledge-graph-based methods such as PGPR (reinforcement learning), CAFE (neuro-symbolic), and PEARLM (generative) perform on educational datasets including COCO, MOOPer, and MOOCube to generate explainable course recommendations [7, 13]. Table 1 presents recommendation utility and path-based explainability metrics for these knowledge-graph-based approaches. These models demonstrate varying degrees of generalizability within education. Among them, PEARLM consistently achieves strong performance, regardless of data characteristics. PEARLM also stands out in explanation type diversity, offering broader and more educationally relevant explanations through a variety of path types. While effective in terms of explainability, this model relies on correlational patterns derived from observed intermediate paths, such as shared course topics, instructors, or institutions-based on user behavior [14]. These paths support interpretability but lack validation of educational dependencies.

Considerations on Causal Inference Integration. To instill causal reasoning, we propose enhancing knowledge-graph-based models with causal constraints that reflect pedagogical logic. For example, a path like: learner $\rightarrow$ enrolled_in $\rightarrow$ Course A $\rightarrow$ prerequisite_for $\rightarrow$ Course B $\rightarrow$ covers $\rightarrow$ Concept C can be validated using causal constraints taken from curriculum data, learner performance records, or instructional design frameworks that represent prerequisite relationships [15]. These constraints allow the system to generate explanations that reflect the causal nature of the recommendation. Rather than stating a generic path-based explanation *Course B is recommended because you took A* the model can generate a more meaningful explanation, *you are recommended to take course B because mastering A causally improves performance in B*. Causal masks can be integrated into the recommendation pipeline, ensuring that Course B is suggested if Course A has been completed, rather than relying on observed past interaction. Consider a scenario where a learner has completed course "Linear Algebra." A correlational model may recommend "Machine Learning" based on frequent co-enrollment patterns. A causal model, however, would evaluate whether the learner took a course on "Probability Theory”, a prerequisite for success in machine learning. By incorporating this causal constraint, the model avoids early recommendations and instead suggests "Probability Theory".

4. Open Challenges and Future Research Directions

Moving from correlational to causal reasoning supports educational decision-making, but also has its own challenges. The primary limitation is the data, as observational data lacks explicit causal annotations. Effective causal inference demands comprehensive learner profiles that include learning objectives and instructional prerequisites, among others. From a modeling perspective, causal reasoning will increase computational complexity. There is also a trade-off between flexibility and faithfulness: causal constraints may improve the validity of recommendations but reduce diversity or novelty, particularly in sparsely-connected data scenarios. Additionally, existing methods are typically evaluated using offline metrics, which do not capture the actual needs in educational recommendations. The lack of user studies further limits the understanding of whether systems support learners’ educational goals. The integration of causal reasoning has the potential to significantly transform educational recommender systems. It offers a strategy that guides learners through structured learning paths. This could have significant implications for instructional design, curriculum development, and learner autonomy. By enabling systems to answer “what if” questions and counterfactual reasoning, causal models empower learners to understand not only what to study next, but why it matters for their long-term learning goals. In summary, we encourage future research to bridge causal inference, language modeling, and educational paradigms, particularly by tackling the open question of how to operationalize causal

Table 1

Performance comparison in terms of recommendation utility (NDCG: Normalized Discounted Cumulative Gain, MRR: Mean Reciprocal Rank) and path-based explainability (LIR: Linking Interaction Recency, LID: Linking Interaction Diversity, SED: Shared Entity Diversity) across three educational datasets. Best results per dataset and metric are in bold; second-best are underlined.

Method	COCO					MOOCube					MOOPer				
	NDCG	MRR	LIR	LID	SED	NDCG	MRR	LIR	LID	SED	NDCG	MRR	LIR	LID	SED
PGPR	0.03	0.03	0.46	**0.68**	**0.81**	**0.11**	0.08	**0.41**	**0.67**	**0.99**	0.32	**0.45**	0.43	**0.66**	**0.99**
CAFE	0.02	0.01	**0.48**	0.04	0.21	0.05	0.04	0.37	0.04	0.22	**0.45**	0.37	**0.44**	0.29	0.29
PLM	**0.04**	**0.04**	0.09	0.39	0.45	0.10	**0.09**	0.07	0.31	0.28	0.28	0.26	0.26	0.59	0.55
PEARLM	0.03	0.02	0.40	0.25	0.40	0.09	0.08	0.32	0.44	0.69	0.31	0.29	0.36	0.60	0.78

reasoning within real educational environments.

Declaration on Generative AI

During the preparation of this work, the author did not use any AI tool.

References

[1] D. Dessì, G. Fenu, M. Marras, D. Reforgiato Recupero, Leveraging cognitive computing for multi-class classification of e-learning videos, in: Proc. of the European Semantic Web Conference, Springer, 2017, pp. 21–25.

[2] D. Dessí, M. Dragoni, G. Fenu, M. Marras, D. Reforgiato Recupero, Deep learning adaptation with word embeddings for sentiment analysis on online course reviews, in: Deep learning-based Approaches for Sentiment Analysis, Springer, 2020, pp. 57–83.

[3] Y. Ge, S. Liu, Z. Fu, J. Tan, Z. Li, S. Xu, Y. Li, Y. Xian, Y. Zhang, A survey on trustworthy recommender systems, Trans. Recomm. Syst. 3 (2025) 13:1–13:68. doi:10.1145/3652891.

[4] J. Stray, A. Halevy, P. Assar, D. Hadfield-Menell, C. Boutilier, A. Ashar, C. Bakalar, L. Beattie, M. Ekstrand, C. Leibowicz, et al., Building human values into recommender systems: An interdisciplinary synthesis, ACM Transactions on Recommender Systems 2 (2024) 1–57.

[5] B. Ma, T. Yang, B. Ren, A survey on explainable course recommendation systems, in: International Conference on Human-Computer Interaction, Springer, 2024, pp. 273–287.

[6] Z. Li, X. Liu, X. Wang, P. Liu, Y. Shen, Transo: a knowledge-driven representation learning method with ontology information constraints, World Wide Web 26 (2023) 297–319.

[7] N. Afreen, G. Balloccu, L. Boratto, G. Fenu, F. M. Malloci, M. Marras, A. G. Martis, Can path-based explainable recommendation methods based on knowledge graphs generalize for personalized education?, in: Proc. of the 33rd ACM Conference on User Modeling, Adaptation and Personalization, 2025, pp. 273–278.

[8] Y. Lin, F. Lin, W. Zeng, J. Xiahou, L. Li, P. Wu, Y. Liu, C. Miao, Hierarchical reinforcement learning with dynamic recurrent mechanism for course recommendation, Knowledge-Based Systems 244 (2022) 108546.

[9] Y. Xian, Z. Fu, S. Muthukrishnan, G. De Melo, Y. Zhang, Reinforcement knowledge graph reasoning for explainable recommendation, in: Proc. of the ACM SIGIR Conference on Research and Development in Information Retrieval, 2019, pp. 285–294.

[10] A. Bakker, E. Angerer, W. R. Penuel, S. F. Akkerman, Causal reasoning about education: What is it and what should it be?, in: The Routledge Handbook of Causality and Causal Methods, Routledge, 2024, pp. 671–682.

[11] M. C. Urdaneta-Ponte, A. Mendez-Zorrilla, I. Oleagordia-Ruiz, Recommendation systems for education: Systematic review, Electronics 10 (2021) 1611.

[12] L. Yang, V. Shirvaikar, O. Clivio, F. Falck, A critical review of causal reasoning benchmarks for large language models, arXiv preprint arXiv:2407.08029 (2024).

[13] G. Balloccu, L. Boratto, C. Cancedda, G. Fenu, M. Marras, Knowledge is power, understanding is impact: Utility and beyond goals, explanation quality, and fairness in path reasoning recommendation, in: Proc. of the European Conference on Information Retrieval, Springer, 2023, pp. 3–19.

[14] N. Afreen, G. Balloccu, L. Boratto, G. Fenu, F. M. Malloci, M. Marras, A. G. Martis, Learner-centered ontology for explainable educational recommendation, in: Adj. Proc. of the 32nd ACM Conference on User Modeling, Adaptation and Personalization, 2024, pp. 567–575.

[15] J. Wang, Y. Le, B. Chang, Y. Wang, E. H. Chi, M. Chen, Learning to augment for casual user recommendation, in: Proc. of the ACM Web Conference, 2022, pp. 2183–2194.