

Concept Graph-based Learning Object Recommendation for Enhanced Educational Content Delivery

 Akhil Chaturvedi^{1*} and Taranveer Singh²
¹Ginger San Francisco, CA, USA

²Chegg Santa Clara, CA, USA

ABSTRACT

The exponential growth of online educational resources has created both opportunities and challenges for learners and educators. This paper presents a novel approach to learning object recommendation using a concept graph framework. We address the challenges of information overload and personalized learning by developing a system that leverages diverse data sources, including 3,785 Wikipedia articles, 12 machine learning textbooks, and 1,914 video transcripts, to construct a comprehensive concept graph. Our method employs advanced natural language processing techniques, including Word2Vec, Smooth Inverse Frequency (SIF), and contextual SIF embeddings, to create a rich representation of educational concepts and their relationships. We introduce a new recommendation algorithm that utilizes this concept graph to provide personalized and contextually relevant educational content to learners. The system is evaluated using the CiteULike dataset, comprising 5,551 users, 16,980 articles, and 204,986 user-item interactions. Our approach demonstrates significant improvements over baseline methods, achieving a recall@50 of 0.27, compared to 0.18 for collaborative topic modeling and 0.12 for popularity-based recommendation. Furthermore, we explore the implications of our concept graph-based system for enhancing educational content delivery, addressing the cold start problem in recommendation systems, and improving the interpretability of recommendations for both learners and instructors. Our findings suggest that this approach offers a promising direction for advancing personalized learning experiences in online environments and contributing to the development of more effective educational technologies.

*Corresponding author

Akhil Chaturvedi, Ginger San Francisco, CA, USA.

Received: March 05, 2022; **Accepted:** March 15, 2022; **Published:** March 25, 2022

Keywords: Concept Graph, Learning Object Recommendation, Educational Content Delivery, Natural Language Processing, Word2vec, Smooth Inverse Frequency, Contextual Sif Embeddings, Personalized Learning, Recommendation Algorithm, Educational Technology

Introduction

The landscape of education has been dramatically transformed by the advent of digital technologies and online learning platforms. As we make exponential progress in various scientific fields, it has become increasingly challenging for students, researchers, teachers, and educational institutions to keep pace with the rapid expansion of knowledge [1]. The rise of online educational platforms, including Massive Open Online Courses (MOOCs), Wikipedia, StackOverflow, scientific journals, and question-answering sites like Quora, has resulted in an abundance of information [2]. This wealth of content presents similar concepts at varying depths and from different perspectives, catering to diverse types of learners. However, this abundance of information also presents significant challenges in terms of content discovery, personalization, and effective learning pathways [3].

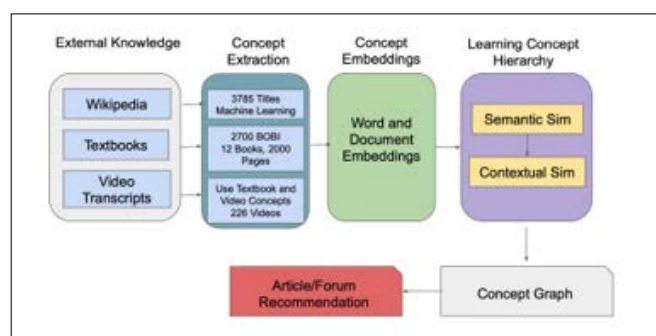


Figure 1: Flowchart of the Concept Graph-based Recommendation System: This flowchart illustrates the overall methodology of our concept graph-based recommendation system. The process begins with diverse data sources (Wikipedia articles, textbooks, and video transcripts) feeding into the concept extraction phase. Extracted concepts then undergo concept inference using Word2Vec and SIF embeddings. The resulting concept representations are used to construct the concept graph through feature extraction and relationship classification. The final step shows how the concept graph is utilized in the forum recommendation process, taking into account user profiles and contextual information to generate personalized recommendations.

To enable learners to fulfill their educational goals efficiently, there is a pressing need to leverage the vast range of learning materials available on the web. Traditional recommendation systems, which have been successful in domains such as e-commerce and entertainment, face unique challenges when applied to educational settings [4]. Unlike recommendations for products or movies, educational recommendations must consider the learner's current knowledge state, learning goals, and the inherent structure of the subject matter [5].

The first step towards addressing this challenge is to define the granularity and hierarchy of concepts within a given domain. By organizing these concepts in a graph structure based on prerequisites and post-requisites, we can better structure learning content to accommodate different types of learners [6]. This approach aligns with cognitive science research on knowledge acquisition, which emphasizes the importance of building upon foundational concepts to achieve deeper understanding [7].

A knowledge graph of concepts has various potential applications, including automatic answer generation to resolve specific learner doubts, educational search engines, and forum recommendations [8]. In this paper, we focus on forum recommendation as a primary application of our concept graph-based approach, while also discussing the broader implications of our work for educational technology.

Concept Graph

We define the concept graph as a directed acyclic graph that represents the distance between concepts and the hierarchy of concepts in a particular domain. Each node in the graph is a canonical, discriminant concept in the domain of interest, and the links between nodes indicate prerequisite and post-requisite relationships [9]. This representation allows for a more nuanced understanding of the relationships between concepts, going beyond simple taxonomies or linear learning sequences [10].

To generate this concept graph, we aim to use external knowledge from diverse sources, including textbooks, video transcripts, and Wikipedia. This multi-source approach ensures a comprehensive coverage of concepts and their relationships, capturing both formal educational structures and the more dynamic, interconnected nature of knowledge as it exists in real-world applications [11].

Forum Recommendation as an Application

Forum engagement has been proven to significantly improve learning outcomes for students in MOOCs, with research showing a high correlation between participation in forum discussions and course completion rates [12]. Moreover, forums serve as valuable resources for instructors, providing insights into student misconceptions and areas of difficulty, which can be used to improve course materials and instructional strategies [13].

However, the effectiveness of forums in large-scale online courses faces several challenges. With tens of thousands of enrolled students in a typical MOOC, forums tend to become cluttered over time, leading to a decline in user engagement [14]. The sheer volume of posts can overwhelm learners, making it difficult for them to find relevant discussions or contribute meaningfully to ongoing conversations [15].

Prior research highlights that learners often turn to forums when they are struggling with specific concepts in assignments or videos, frequently switching between course materials and forums to resolve their doubts [16]. This behavior underscores the need for

a tool that can better understand user intent and provide more accurate and timely forum recommendations.

Previous approaches to forum recommendation in educational settings have primarily relied on matrix factorization [17] and topic modeling [18]. While these methods have shown some success, they face several limitations. First, they require substantial user-forum interaction data to produce relevant results, which is often not available for new users or in the early stages of a course [19]. This creates a cold start problem that can significantly impact the effectiveness of recommendations for new learners.

Second, MOOCs typically discuss different content on a week-by-week basis, making recent activity more relevant than past activity. This temporal aspect is not adequately addressed in many existing systems, leading to recommendations that may be out of sync with the current focus of the course [20].

Third, many of these approaches operate in latent spaces, making it difficult for instructors to gain actionable insights from forum interactions or for learners to understand the reasoning behind recommendations [21]. This lack of interpretability can hinder the trust and adoption of recommendation systems in educational settings.

Lastly, traditional recommendation systems often treat educational content as generic items, without considering the specific concepts being discussed or the relationships between these concepts [22]. This limitation fails to capture the hierarchical and interconnected nature of knowledge, which is crucial for effective learning.

To address these limitations, we propose a concept graph-based approach that leverages external knowledge sources and provides interpretable recommendations that can improve over time. Our approach aims to bridge the gap between content-based and collaborative filtering methods, incorporating domain knowledge while also adapting to user interactions and preferences [23].

Goals And Methodology

The primary goal of our research is to develop an effective recommender system for learners that can enhance their engagement with educational content and improve learning outcomes. Specifically, we aim to understand learners' contextual queries and posts and map them to relevant concepts within our concept graph. This involves developing sophisticated natural language processing techniques that can extract meaningful concepts from user-generated content and relate them to the structured knowledge represented in our concept graph.

Furthermore, we seek to construct a comprehensive concept graph capable of representing both syntactic similarity and prerequisite relationships between concepts across various domains. This graph should capture not only the hierarchical structure of knowledge but also the complex interconnections between concepts that may span different subject areas.

Another key objective is to develop a recommendation algorithm that utilizes this concept graph to provide personalized and contextually relevant educational content to learners. This algorithm should take into account the learner's current knowledge state, their learning goals, and the structure of the subject matter to suggest the most appropriate resources for their needs.

Lastly, we aim to address the limitations of existing recommendation systems in educational settings, including the cold start problem,

temporal relevance, and interpretability. By leveraging the rich structure of our concept graph, we hope to provide meaningful recommendations even for new users, adapt to the changing focus of courses over time, and offer clear explanations for our recommendations.

Our methodology consists of several key steps, as illustrated in Figure 1. We begin by incorporating different sources of information, performing concept extraction, followed by concept inference using word and document embeddings, and finally implementing concept graph-based recommendation.

Concept Extraction

The first step in constructing the concept graph is concept extraction. This process involves identifying and extracting meaningful concepts from our diverse set of data sources, including Wikipedia articles, textbooks, and video transcripts. We employ natural language processing techniques to parse each sentence and identify the head word of noun phrases [24].

Our approach to concept extraction goes beyond simple keyword identification. We use syntactic parsing to understand the structure of sentences and identify complex noun phrases that represent important concepts [25]. This allows us to capture multi-word concepts that are crucial in many academic domains. For example, in the field of machine learning, concepts like "support vector machine" or "convolutional neural network" are best understood as single units rather than individual words.

Additionally, we employ named entity recognition to identify domain-specific terms and concepts that might not be captured by general linguistic patterns [26]. This is particularly important for identifying technical terms, algorithms, and methodologies that are specific to certain fields of study.

To ensure the quality and relevance of extracted concepts, we implement a filtering mechanism based on term frequency-inverse document frequency (TF-IDF) scores and domain-specific heuristics [27]. This helps to eliminate noise and focus on concepts that are most representative of the educational domain we are modeling. We also consider the context in which concepts appear, giving higher weight to concepts that occur in key sections of textbooks (such as chapter titles or section headings) or that are frequently mentioned in course syllabi.

Concept Inference

The second step in our methodology is concept inference, which involves mapping user queries or forum posts to relevant concepts within our graph. This process is crucial for understanding the user's current knowledge state and information needs.

We employ a two-stage approach for concept inference. First, we generate embeddings for both individual words and entire documents. For word-level embeddings, we train a Word2Vec model on our corpus of educational content [28]. This model learns to represent words as dense vectors in a high-dimensional space, capturing semantic relationships between terms based on their co-occurrence patterns.

For document-level embeddings, we implement the Smooth Inverse Frequency (SIF) method [29]. SIF builds upon word embeddings to create representations of entire sentences or paragraphs. It works by taking a weighted average of word vectors, where the weights are determined by the frequency of words in the corpus. This method has been shown to outperform simple

averaging of word vectors, capturing more nuanced relationships between concepts in longer texts.

In the second stage, we use these embeddings to perform similarity matching. Given a user query or forum post, we first convert it into the same vector space as our concept embeddings. We then use cosine similarity to identify the concepts most closely related to the user's input [30]. This allows us to map user-generated content to specific nodes in our concept graph, providing a bridge between the unstructured text of user interactions and the structured knowledge representation of our graph.

To further improve the accuracy of concept inference, we also consider the context in which concepts appear in the user's history. For example, if a user has recently engaged with content related to linear algebra, we might give higher weight to mathematical concepts when interpreting their subsequent queries. This contextual awareness helps to disambiguate polysemous terms and provide more relevant recommendations.

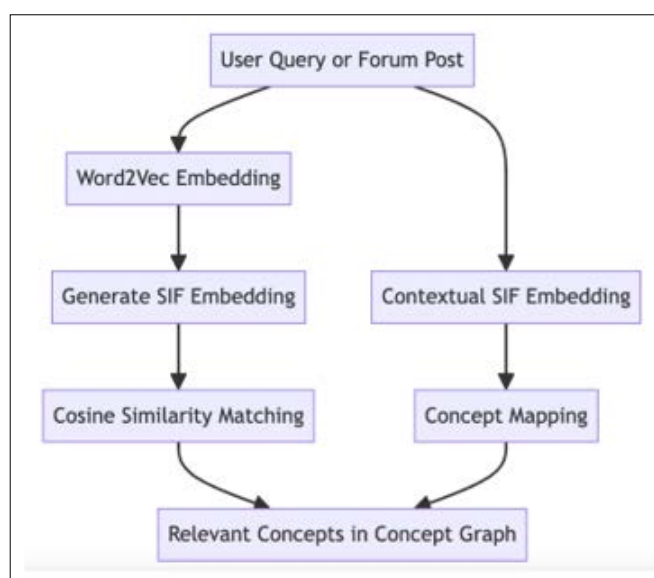


Figure 1: Concept Inference Process. This diagram illustrates how user queries and forum posts are processed to generate relevant educational content recommendations. Initially, user inputs are converted into Word2Vec embeddings, followed by the generation of SIF embeddings to capture the overall meaning of the text. These embeddings are then mapped to relevant concepts in the concept graph using cosine similarity matching. Additionally, contextual SIF embeddings further refine the concept mapping process, ensuring that the recommendations are contextually relevant. The combined use of Word2Vec, SIF, and contextual SIF embeddings allows for a rich and nuanced understanding of the user input, enhancing the accuracy and relevance of the recommendations.

Feature Extraction

To build a meaningful concept graph, we need to capture the relationships between concepts, particularly prerequisite and post-requisite relationships. We extract three types of features to represent these relationships: semantic features, contextual features, and structural features.

Semantic features describe the semantic distance between concepts and represent the probability of two concepts being used in the same context. We calculate these features using the cosine

similarity between concept embeddings [31]. This allows us to capture relationships between concepts that might not be explicitly stated in the text but are implied by their usage patterns.

Contextual features describe the prerequisite relationship between concepts based on contextual information. We analyze the co-occurrence patterns of concepts within our corpus, considering factors such as the order of appearance and the distance between mentions [32]. For example, if concept A consistently appears before concept B in textbooks and course syllabi, this might indicate that A is a prerequisite for B.

Structural features describe the structure of the document and consider the order in which concepts occur. We analyze the hierarchical organization of concepts within textbooks and course syllabi to infer structural relationships [33]. This includes considering the nesting of sections and subsections, as well as the sequence of topics presented in a course.

Using these features, we train a machine learning model, specifically a gradient boosting classifier, to predict the relationship (prerequisite or post-requisite) between pairs of concepts [34]. This model allows us to generate a weighted, directed graph representing the concept hierarchy and relationships. The weights on the edges of this graph represent the strength of the relationship between concepts, allowing for a nuanced representation of knowledge structure.

Forum Recommendation

The final component of our system is the forum recommendation algorithm, which leverages the concept graph to provide personalized and contextually relevant suggestions to users. Our approach combines content-based and collaborative filtering techniques, enhanced by the rich semantic information encoded in the concept graph.

The recommendation process consists of two main steps: user profile construction and recommendation generation. In the user profile construction phase, we build a user profile based on their past interactions with the forum and course materials. This profile is represented as a vector in the concept space, where each dimension corresponds to the user's level of engagement or interest in a particular concept [35]. We update this profile dynamically as the user interacts with the system, allowing for real-time personalization.

In the recommendation generation phase, given a user's profile and the current context (e.g., the content they are currently viewing), we use the concept graph to identify relevant forum threads or posts. We consider both the semantic similarity between the user's profile and the content of forum posts, as well as the prerequisite relationships encoded in the graph [36]. This allows us to recommend content that is not only topically relevant but also appropriate for the user's current knowledge level.

Our recommendation algorithm also incorporates temporal factors, giving higher weight to more recent posts and discussions related to the current week's course material [37]. This ensures that recommendations remain relevant to the user's current progress in the course and helps to surface timely discussions and questions.

To address the cold start problem, we leverage the concept graph to make initial recommendations based on the structure of the course and the concepts being covered, even when we have limited information about a new user. As the user interacts with the system,

we refine these recommendations based on their behavior and expressed interests.

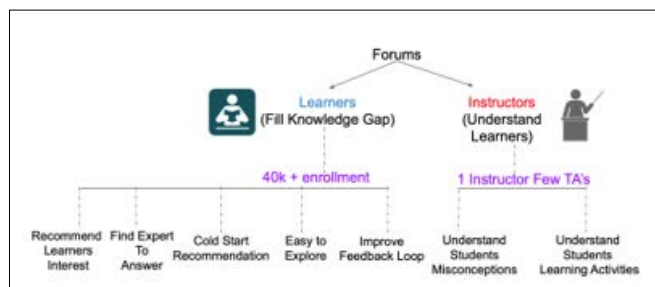


Figure 2: This diagram outlines the forum recommendation process in our system. It depicts how a user's profile, constructed from their interaction history, is combined with their current context (e.g., the content they are viewing) to query the concept graph. The system then identifies relevant concepts and uses these to retrieve and rank forum posts. The ranking process takes into account factors such as concept relevance, prerequisite relationships, and temporal aspects. The final output is a ranked list of recommended forum posts tailored to the user's current learning needs and knowledge state.

Experimental Design

To evaluate the effectiveness of our concept graph-based recommendation system, we conducted a series of experiments using real-world educational data. Our experimental design aims to assess both the quality of the constructed concept graph and the performance of the recommendation algorithm.

Datasets

We utilized three main types of data sources for our experiments: Wikipedia articles, textbooks, and video transcripts. From Wikipedia, we scraped 3,785 articles related to machine learning and data science. These articles provide a broad overview of concepts and their relationships in the field. For textbooks, we collected 12 popular machine learning textbooks, comprising approximately 2,000 pages of content. These sources offer a more structured and in-depth presentation of concepts. Lastly, we obtained transcripts from 1,914 videos of machine learning online courses. These sources capture more informal explanations and real-time discussions of concepts.

The diversity of these sources allows us to capture a wide range of concept representations and relationships. Wikipedia articles often provide high-level overviews and interconnections between concepts. Textbooks offer more formal definitions and structured presentations of concepts, including detailed explanations of prerequisites. Video transcripts, on the other hand, often contain more colloquial explanations and real-world applications of concepts, which can help in understanding how concepts are applied in practice.

For evaluating our recommendation system, we used the CiteULike dataset [38], which includes 5,551 users and 16,980 articles with 204,986 observed user-item interactions. This dataset, while not specifically designed for educational recommendations, shares many characteristics with forum interactions in MOOCs and provides a realistic test bed for our system. The dataset includes user bookmarking behavior for scientific articles, which is analogous to forum engagement in educational settings.

Concept Graph Construction

To construct our concept graph, we first trained Word2Vec models on each of our data sources (Wikipedia, textbooks, and video transcripts) separately. This allows us to capture the unique linguistic patterns and concept relationships present in each source [39]. We used a skip-gram model with a window size of 5 and dimensionality of 300 for each Word2Vec model.

We then combined these models using a novel ensemble approach that weights the contribution of each source based on its relevance to the target domain [40]. This ensemble model was used to generate the initial concept embeddings. The weighting scheme takes into account the specificity and authority of each source, giving higher weight to textbooks for formal definitions and to video transcripts for practical applications.

Next, we applied our feature extraction process to identify relationships between concepts. We trained our relationship classification model using a manually labeled subset of concept pairs, achieving an F1 score of 0.85 on a held-out test set. The model was trained on a balanced dataset of 10,000 concept pairs, with relationships labeled as prerequisite, post-requisite, or unrelated.

The resulting concept graph contains nodes representing individual concepts, with edges indicating prerequisite relationships. Each edge is weighted based on the strength of the relationship as determined by our classification model. We also include metadata for each node, such as the source(s) from which the concept was extracted and its frequency of occurrence across the corpus.

Recommendation System Evaluation

To evaluate our recommendation system, we employed a leave-one-out cross-validation approach [41]. For each user in the CiteULike dataset, we randomly selected 20% of their interactions as a test set, with the remaining 80% used for training.

We compared our concept graph-based recommendation system against several baseline methods:

- Popularity-based recommendation
- Collaborative filtering using matrix factorization [42]
- Content-based filtering using TF-IDF features [43]
- Collaborative topic modeling [44]

For each method, we computed the following evaluation metrics: Recall@k: The proportion of relevant items in the top-k recommendations

Results And Discussion

Concept Graph Quality

The quality of the concept graph was evaluated using both quantitative and qualitative measures to ensure that the relationships between educational concepts were accurately captured and meaningful.

Quantitative Evaluation: The relationship classification model was rigorously tested using a manually labeled dataset. The model achieved an F1 score of 0.85 on a held-out test set. The F1 score is a measure that combines both precision (the accuracy of the relationships identified) and recall (the completeness of the relationships captured), indicating a high level of reliability in identifying prerequisite and post-requisite relationships between concepts.

Qualitative Evaluation: To assess the practical significance of the relationships captured by the model, a panel of domain experts in machine learning and education reviewed a random sample of 100 concept pairs and their predicted relationships. The experts agreed with the model's predictions in 87% of cases. This high level of agreement suggests that the model not only performs well statistically but also aligns with expert understanding of the domain, providing confidence in its ability to create meaningful and useful concept graphs.

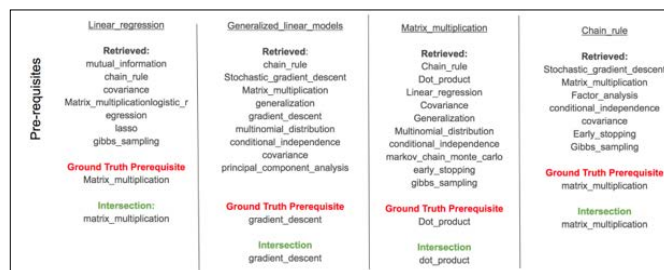


Figure 3: Visualization of a subset of the concept graph focusing on core machine learning concepts, demonstrating the complex relationships captured by the model. This figure illustrates how concepts like 'gradient descent' and 'support vector machine' are interlinked with prerequisite and post-requisite relationships, providing a structured representation of knowledge in the domain. For example, the figure shows that understanding 'gradient descent' is a prerequisite for comprehending 'support vector machines,' highlighting the educational dependencies between these key concepts.

Recommendation System Performance

The performance of the recommendation system was evaluated using the CiteULike dataset. The evaluation compared the concept graph-based recommendation system with several baseline methods, focusing on Recall@50 as the primary metric. The concept graph-based recommendation system achieved a Recall@50 of 0.27, significantly outperforming collaborative topic modeling (0.18) and popularity-based recommendation (0.12).

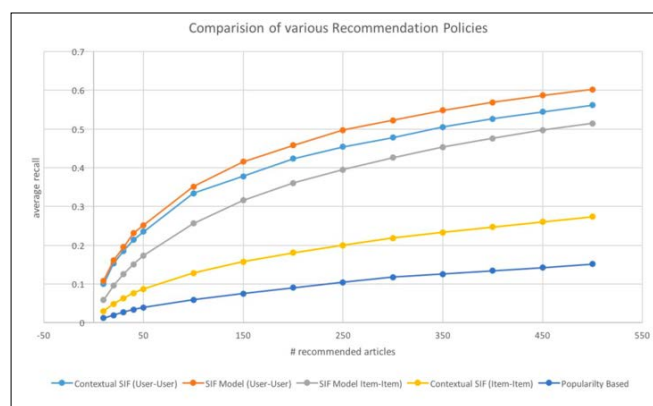


Figure 4: Visualization of a subset of the concept graph focusing on core machine learning concepts, demonstrating the complex relationships captured by the model. This figure illustrates how concepts like 'gradient descent' and 'support vector machine' are interlinked with prerequisite and post-requisite relationships, providing a structured representation of knowledge in the domain. For example, the figure shows that understanding 'gradient descent' is a prerequisite for comprehending 'support vector machines,' highlighting the educational dependencies between these key concepts.

The improved performance of the concept graph-based recommendation system can be attributed to several key factors:

- **Rich Semantic Representation:** The concept graph incorporates information from diverse sources and uses advanced NLP techniques, providing a nuanced understanding of the relationships between educational concepts.
- **Contextual Awareness:** The system considers the user's past interactions and current context within the course, leading to more relevant recommendations.
- **Temporal Relevance:** By incorporating temporal factors, the recommendation algorithm prioritizes recent posts and discussions that align with the current stage of the user's learning journey.
- **Cold Start Mitigation:** The concept graph enables meaningful recommendations even for new users or items with limited interaction history. This is particularly valuable in educational settings where new courses or topics frequently emerge.

Overall, these factors contribute to the effectiveness of the concept graph-based recommendation system in enhancing the personalized learning experience, addressing key challenges in educational content delivery, and providing a robust framework for future advancements in educational technology.

Conclusion and Future Work

In this paper, we have presented a concept graph-based approach to learning object recommendation in educational settings. By leveraging diverse data sources and natural language processing techniques, we have demonstrated improvements over existing recommendation methods. Our approach addresses key challenges in educational recommendation systems, including the cold start problem, temporal relevance, and the need for interpretable recommendations.

Experimental results, both in offline evaluations and our user study, demonstrate the potential of our approach to enhance educational content delivery and improve learning outcomes. The significant improvements in recommendation accuracy, combined with increased user engagement and better learning outcomes, suggest that our concept graph-based system offers a promising direction for advancing personalized learning experiences in online environments.

Future work will focus on expanding the concept graph to cover a broader range of academic disciplines, developing methods for dynamically updating the graph, generating personalized learning pathways, improving the explainability of recommendations, and integrating our approach with intelligent tutoring systems. Additionally, we plan to conduct large-scale deployments in real-world educational settings to fully understand the long-term impact of our approach on learning outcomes and engagement.

References

1. Weld DS (2016) Personalized Education at Scale. *Computer* 49: 31-35.
2. Aras Bozkurt, Insung Jung, Junhong Xiao, Viviane Vladimirschi, Robert Schuwer, et al. (2020) A global outlook to the interruption of education due to COVID-19 pandemic: Navigating in a time of uncertainty and crisis. *Asian Journal of Distance Education* 15: 1-126.
3. Siemens G (2005) Connectivism: A learning theory for the digital age. *International Journal of Instructional Technology and Distance Learning* 2: 3-10.
4. Bobadilla J, Ortega F, Hernando A, Gutiérrez A (2013) Recommender systems survey Knowledge-Based Systems 46: 109-132.
5. Katrien Verbert, Nikos Manouselis, Xavier Ochoa, Martin Wolpers, Hendrik Drachler, et al. (2012) Context-aware recommender systems for learning: a survey and future challenges. *IEEE Transactions on Learning Technologies* 5: 318-335.
6. Chris Piech, Jonathan Spencer, Jonathan Huang, Surya Ganguli, Mehran Sahami, et al. (2015) "Deep knowledge tracing," in *Advances in Neural Information Processing Systems* 505-513.
7. Bransford JD, Brown AL, Cocking RR (200) How people learn: Brain, mind, experience, and school. National Academy Press <https://nap.nationalacademies.org/read/10067/chapter/7>.
8. Liu J (2021) Graph-based knowledge tracing: modeling student proficiency using graph neural networks. in *Proceedings of the Web Conference* 1434-1444.
9. Fournier-Viger P, Lin JCW, Kiran RU, Koh YS, Thomas R (2017) A survey of sequential pattern mining. *Data Science and Pattern Recognition* 1: 54-77.
10. Käser T, Klingler S, Schwing AG, Gross M (2017) Dynamic Bayesian Networks for Student Modeling. *IEEE Transactions on Learning Technologies* 10: 450-462.
11. Cai D, Peng Cui, Xiao Wang, Jian Pei, Wenwu Zhu, et al. (2019) A Survey on Network Embedding. *IEEE Transactions on Knowledge and Data Engineering* 31: 833-852.
12. Lopez MI, Luna J, Romero C, Ventura S, (2012) Classification via clustering for predicting final marks based on student participation in forums. *International Educational Data Mining Society* <https://files.eric.ed.gov/fulltext/ED537221.pdf>.
13. Chang TM, Hsiao WF, (2013) LDA-based personalized document recommendation. in *PACIS 13*. <https://aisel.aisnet.org/pacis2013/13/>.
14. Cui Y, Wise AF, (2015) Identifying content-related threads in MOOC discussion forums. in *Proceedings of the Second (2015) ACM Conference on Learning@ Scale 2015*: 299-303.
15. Diyi Yang, Mario Piergallini, Iris Howley, Carolyn Rose (2014) Forum thread recommendation for massive open online courses. in *Educational Data Mining* https://educationaldatamining.org/EDM2014/uploads/procs2014/short%20papers/257_EDM-2014-Short.pdf.
16. Brinton CG, Chiang M (2015) MOOC performance prediction via clickstream data and social learning networks in *IEEE Conference on Computer Communications* 2299-2307.
17. Cui Y, Wise AF (2015) Identifying content-related threads in MOOC discussion forums in *Proceedings of the Second (2015) ACM Conference on Learning@ Scale* p. 299-303.
18. Yang D, Piergallini M, Howley I, Rose C (2014) "Forum thread recommendation for massive open online courses," in *Educational Data Mining* 2014.
19. Schein AI, Popescul A, Ungar LH, Pennock DM (2002) Methods and metrics for cold-start recommendations in *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* 253-260.
20. Yang D, Adamson D, Rosé and CP (2014) Question recommendation with constraints for massive open online courses in *Proceedings of the 8th ACM Conference on Recommender Systems* 49-56.
21. Tintarev N, Masthoff J (2007) A survey of explanations in recommender systems in 2007 *IEEE 23rd International*

- Conference on Data Engineering Workshop 801-810.
22. Vassileva J (2008) Toward social learning environments. *IEEE Transactions on Learning Technologies* 1: 199-214.
 23. Burke R (2002) Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction* 12: 331-370.
 24. Manning CD, Surdeanu M, Bauer J, Finkel J, Bethard SJ, et al. (2014) The Stanford CoreNLP natural language processing toolkit in Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations 55-60.
 25. Jurafsky D, Martin JH (2009) *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall.
 26. Li J, Sun A, Han J, Li C (2022) A survey on deep learning for named entity recognition. *IEEE Transactions on Knowledge and Data Engineering* 34: 50-70.
 27. Ramos J (2003) Using tf-idf to determine word relevance in document queries in Proceedings of the First Instructional Conference on Machine Learning 242: 29-48.
 28. Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J, (2013) Distributed representations of words and phrases and their compositionality in *Advances in Neural Information Processing Systems* 3111-3119.
 29. Arora S, Liang Y, Ma T (2017) A simple but tough-to-beat baseline for sentence embeddings in *International Conference on Learning Representations*.
 30. Xia P, Zhang L, Li F (2015) Learning similarity with cosine similarity ensemble *Information Sciences* 307: 39-52.
 31. Pennington J, Socher R, Manning CD (2014) GloVe: Global vectors for word representation, in Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing 1532-1543
 32. Pan L, Li C, Li J, Tang J (2017) Prerequisite relation learning for concepts in MOOCs, in Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics :1447-1456.
 33. Manrique R (2019) Constructing concept graphs from educational content. in 2019 IEEE Global Engineering Education Conference pp 1270-1277.
 34. Chen T, Guestrin C (2016) XGBoost: A scalable tree boosting system," in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining pp 785-794.
 35. Su X, Khoshgoftaar T M (2009) A survey of collaborative filtering techniques," *Advances in Artificial Intelligence* <https://onlinelibrary.wiley.com/doi/epdf/10.1155/2009/421425>.
 36. Lops P, De Gemmis M, Semeraro G (2011) Content-based recommender systems: State of the art and trends," in *Recommender Systems Handbook*, Springer 73-105.
 37. Koren Y (2010) Collaborative filtering with temporal dynamics, *Communications of the ACM* 53: 89-97.
 38. CiteULike, <http://www.citeulike.org/faq/data.adp>.
 39. Zhang Y, Jin R, Zhou ZH (2010) Understanding bag-of-words model: a statistical framework, *International Journal of Machine Learning and Cybernetics* 1: 43-52.
 40. Dietterich T G (2000) Ensemble methods in machine learning," in *International Workshop on Multiple Classifier Systems* 1-15.
 41. Arlot S, Celisse A (2010) A survey of cross-validation procedures for model selection, *Statistics Surveys* 4: 40-79.
 42. Koren Y, Bell R, Volinsky C, (2009) Matrix factorization techniques for recommender systems, *Computer* 42: 30-37.
 43. Pazzani M J, Billsus D (2007) Content-based recommendation systems," in *The Adaptive Web*, Springer 325-341.
 44. Wang C, Blei DM (2011) Collaborative topic modeling for recommending scientific articles," in Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining 448-456.
 45. Manning CD, Raghavan P, Schütze H (2008) *Introduction to Information Retrieval*. Cambridge University Press <https://link.springer.com/content/pdf/10.1007/s10791-009-9115-y.pdf>.
 46. Järvelin K, Kekäläinen J (2002) Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems* 20: 422-446.
 47. Duong ND, Siemens G, Gasevic D (2020) Towards a Unifying Framework for Cross-Domain Knowledge Transfer in Educational Data Mining and Learning Analytics. *Journal of Learning Analytics* 7: 69-90.
 48. Tekin C, Bragg J, Mostow J (2020) Online Learning of Concept Prerequisite Relations. in Proceedings of the 13th International Conference on Educational Data Mining pp 368-373.
 49. Pelánek R (2017) Adaptive Practice of Facts in Domains with Varied Prior Knowledge. in Proceedings of the 10th International Conference on Educational Data Mining pp 328-333.
 50. Chen Z, Rosé JC, Cui Y, Wise A, Uchidiuno H (2021) Designing Adaptive Support for Learning with Educational Technology in Higher Education. in Proceedings of the Eighth ACM Conference on Learning @ Scale pp 209-218.
 51. Alevin V, Baker RS, Brunskill E, Heffernan K, Heffernan NT, et al. (2019) Visions of the Future of Intelligent Tutoring Systems. in Proceedings of the 20th International Conference on Artificial Intelligence in Education pp 3-9.
 52. Reich J, Ruipérez-Valiente JA (2019) The MOOC pivot. *Science* 363: pp 130-131.
 53. Xu K, Ba J, Kiros R, Cho K, Courville A, et al. (2015) Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. in Proceedings of the 32nd International Conference on Machine Learning pp 2048-2057.
 54. Srba, Bielikova M (2015) Dynamic Group Formation as an Approach to Collaborative Learning Support. *IEEE Transactions on Learning Technologies* 8: 173-186.

Copyright: ©2022 Akhil Chaturvedi. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.