INTERNATIONAL JOURNAL OF INFORMATION TECHNOLOGY RESEARCH STUDIES (IJITRS)
(Open Access, Double-Blind Peer Reviewed Journal)
ISSN Online: 3139-3209

Personalized Course Recommendation on MOOC Platforms: A Survey

Rejina P V

Assistant professor, Co-Operative Arts And Science College, Madayi ,Pazhayangadi, Kannur, India

Article information

Received: 23rd March 2026
Received in revised form: 25th April 2026
Accepted: 27th May 2026
Available online: 30th July 2026

Volume: 2
Issue: 3
DOI: https://doi.org/10.63090/IJITRS/3139.3209.0031

Abstract

Massive Open Online Courses (MOOCs) delivered through platforms such as Coursera, edX, Udacity, FutureLearn, XuetangX and SWAYAM have placed tens of thousands of courses within reach of millions of learners. That abundance has produced a side effect: acute information overload, difficult course selection and persistently high dropout. Recommendation systems have become the principal mechanism for guiding learners toward suitable courses, concepts and learning paths. This paper presents a structured review of the literature on course recommendation for MOOC platforms. A taxonomy is developed that organises prior work into collaborative filtering (memory- and model-based, including matrix factorization and Bayesian personalized ranking), content-based and knowledge-based methods, sequential and session-based models, knowledge-graph-based approaches, deep-learning architectures (neural collaborative filtering, recurrent and attention networks), reinforcement-learning formulations and hybrid designs, together with concept and learning-path recommendation. The benchmark datasets and platforms that underpin reproducible research, including MOOCCube, MOOCCubeX and XuetangX, are summarised. The evaluation metrics used across the field are then compared: ranking measures such as precision, recall, F1@K, NDCG, MAP, hit-rate and MRR alongside learning-centric indicators. The review synthesises the open challenges of cold-start, data sparsity, dropout-aware recommendation, scalability, explainability, fairness and diversity, pedagogical alignment with the learner knowledge state, and privacy. Finally, emerging directions are outlined. These span large-language-model and conversational recommenders, knowledge-tracing-aware systems, federated and privacy-preserving recommendation, and fairness. The survey is intended as a consolidated reference for researchers and practitioners building learner-centred recommendation systems in online education.

Keywords:- MOOC, Course Recommendation, Recommender Systems, Collaborative Filtering, Knowledge Graph; Deep Learning, Online Education, Survey.

I. INTRODUCTION

Massive Open Online Courses (MOOCs) have reshaped access to higher and continuing education over the past decade. Platforms such as Coursera, edX, Udacity, FutureLearn, China's XuetangX and India's SWAYAM collectively host tens of thousands of courses across nearly every discipline. They enrol learners who differ widely in background, goal and prior knowledge [1], [2]. Scale democratises learning, yet it also creates an information-overload problem. A prospective learner faces a catalogue far too large to inspect by hand, and an unsuitable choice, too advanced, too elementary, or poorly aligned with a career objective, contributes to disengagement. MOOCs are well known for low completion; many courses retain only a small fraction of enrolled learners through to certification [3], [4]. Course recommendation is therefore not a convenience feature. It is a central lever for learner success and retention.

Recommendation systems were first studied in e-commerce and media [5], [6], and have since been adapted extensively to the MOOC setting. The educational context imposes distinctive requirements. Recommendations must respect the sequential and prerequisite structure of knowledge, account for the evolving competence of the learner, and ideally advance pedagogical goals rather than mere click-through. Over time the field has moved from classical collaborative filtering and matrix factorization to deep neural, sequence-aware and knowledge-graph-based models, and most recently toward large-language-model and conversational interfaces. This paper reviews that trajectory.

The contribution of this survey is fourfold. First, it organises the literature into a coherent taxonomy of methodological families (Fig. 1). Second, it documents the datasets, platforms and evaluation metrics that support reproducible research. Third, it synthesises the open challenges that recur across studies. Fourth, it outlines forward-looking directions. The remainder of the paper is structured accordingly. Section II frames the MOOC recommendation problem. Section III presents the taxonomy and methods. Section IV surveys datasets and platforms. Section V reviews evaluation metrics. Section VI discusses open challenges. Section VII outlines future directions, and Section VIII concludes. This work is a review only; it reports no new system or experiments of its own.

II. MOOCS AND THE COURSE-RECOMMENDATION PROBLEM

A MOOC platform can be modelled as an interaction environment. A set of learners engages with a set of items, courses, videos, exercises or fine-grained concepts, generating explicit signals such as ratings and enrolments and implicit signals such as clicks, watch time and assessment outcomes. The recommendation task is to predict, for each learner, the items most likely to be useful next. Retail recommendation captures value through a purchase; education does not. A course that is clicked but never completed, or completed without learning gain, is a weak success. The objective therefore couples relevance with learning outcome.

Three structural difficulties characterise the setting. Data are extremely sparse, because most learners interact with only a handful of the available courses, leaving the learner-item matrix almost entirely empty [5], [7]. New learners and newly published courses arrive continuously, producing the cold-start condition in which no interaction history exists [8]. And learner intent is non-stationary: a sequence of enrolments reflects a progression through prerequisite knowledge, so order carries meaning that static models ignore [9], [10]. A generic recommendation pipeline that addresses these difficulties appears in Fig. 2. It comprises data ingestion, feature or embedding learning, candidate generation, ranking, and a feedback loop that returns interaction and learning signals to the model.

Closely related is the problem of dropout, studied both as a prediction task and as a signal for recommendation. Early work framed MOOC dropout as a supervised learning problem over clickstream features [3]. Later studies modelled it at scale to understand its drivers [4]. Because an ill-matched course raises the likelihood of disengagement, dropout prediction and course recommendation are increasingly treated as complementary components of a single retention strategy.

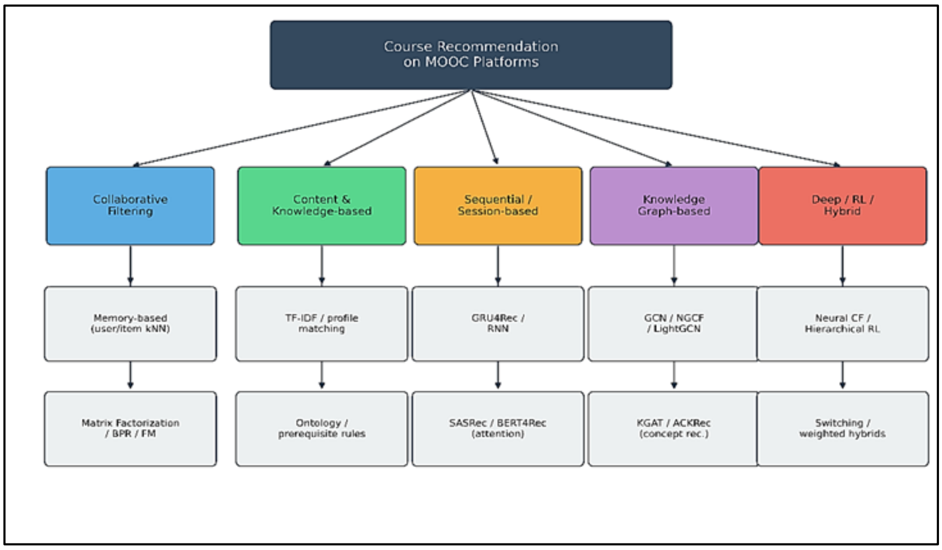

Fig. 1: Taxonomy of recommendation approaches surveyed for MOOC course recommendation.

III. TAXONOMY OF RECOMMENDATION METHODS

This section organises the methodological literature into the families shown in Fig. 1 and compared in Table 1. The families are not mutually exclusive. Many practical systems combine them, and the most recent work blends deep representation learning with knowledge-graph structure and sequence modelling.

A. Collaborative Filtering

Collaborative filtering (CF) recommends items by exploiting patterns of agreement among users. It remains the workhorse of the field. Memory-based CF computes similarities directly between users or between items; the item-based formulation of Sarwar et al. [7] is particularly influential, because item-item similarities are comparatively stable and precomputable. Model-based CF instead learns latent representations. Matrix factorization, popularised for recommendation by Koren et al. [5], decomposes the sparse interaction matrix into low-dimensional user and item factors and underpins a large body of MOOC work. For the implicit-feedback signals typical of MOOCs, Bayesian Personalized Ranking (BPR) [11] reframes learning as a ranking problem over observed versus unobserved items, while Factorization Machines [12] generalise matrix factorization to arbitrary side features. CF is effective and domain-agnostic, yet it degrades under sparsity and cold-start.

B. Content-Based and Knowledge-Based Methods

Content-based methods recommend courses whose descriptive features, syllabus text, topics and skills, resemble those a learner has previously favoured, building an explicit learner profile [13], [14]. They rely on item content rather than co-interaction, so they ease item cold-start and offer transparent explanations. The trade-off is over-specialisation and limited novelty. Knowledge-based methods encode explicit domain constraints, prerequisite relationships, competency requirements or curricular rules, and reason over them to recommend courses or concepts. In the MOOC domain, learning prerequisite relations among concepts is itself an active task [15]. Prerequisite structure is a natural foundation for learning-path recommendation that respects the order in which knowledge should be acquired.

C. Sequential and Session-Based Recommendation

Learning unfolds over time, so sequence-aware models are especially pertinent. Session-based recommendation with recurrent networks, exemplified by GRU4Rec [9], models a user's recent activity as an ordered sequence and predicts the next item. Self-attentive models then improved both accuracy and efficiency. SASRec [16] applies a unidirectional Transformer to interaction sequences, and BERT4Rec [17] uses a bidirectional masked objective. Graph-structured session models such as SR-GNN [10] capture complex item transitions within a session. In MOOCs these models naturally represent the progression of a learner through successive courses or concepts.

D. Deep-Learning Approaches

Deep learning has reshaped recommendation broadly [18]. Neural Collaborative Filtering [19] replaces the inner product of matrix factorization with a learned interaction function, while deep architectures for industrial-scale candidate generation and ranking were demonstrated by Covington et al. [20]. The attention mechanism of the Transformer [21] now permeates sequential and feature-interaction models alike. Within MOOCs, deep models have been applied to resource and course recommendation; hierarchical reinforcement learning over learner profiles is one prominent example, discussed in Section III-F [22]. The strength of deep methods is representational flexibility. Their costs are data hunger, computational expense and reduced interpretability.

E. Knowledge-Graph-Based Recommendation

Knowledge graphs (KGs) inject relational structure, courses, concepts, teachers and prerequisites, into the recommender, improving both accuracy and explainability. Graph convolutional networks [23] and their recommendation-specific descendants, Neural Graph Collaborative Filtering [24] and LightGCN [25], propagate information along the user-item graph. KGAT [26] attends over a unified knowledge graph to model high-order connectivity; a broader survey of graph neural networks in recommendation is given in [28]. In the MOOC setting specifically, ACKRec [27] casts knowledge-concept recommendation as learning over a heterogeneous information network with attentional graph convolution, exploiting the rich entity structure that MOOC repositories expose. KG-based methods suit concept-level and learning-path recommendation particularly well.

F. Reinforcement-Learning Approaches

Reinforcement learning (RL) frames recommendation as sequential decision-making that optimises a long-horizon objective, such as sustained engagement or learning gain, rather than a single next-item prediction. A prominent MOOC example is the hierarchical RL course-recommendation model of Zhang et al. [22]. A high-level agent revises the learner profile by removing noisy historical signals, and a low-level agent issues recommendations; the two are trained jointly on real MOOC data. RL aligns with the delayed, cumulative nature of educational reward. It also raises hard questions of reward design, sample efficiency and safe online exploration.

G. Hybrid Methods and Concept / Learning-Path Recommendation

Hybrid recommenders combine two or more of the above families to offset their individual weaknesses. Burke's taxonomy of weighted, switching, cascade and feature-combination hybrids remains the standard reference [29]. In practice, MOOC systems frequently blend collaborative signal with content or knowledge-graph structure to handle cold-start while retaining personalization. A distinctive educational specialisation is the move from recommending whole courses to recommending fine-grained concepts and ordered learning paths, building on prerequisite-relation learning [15] and concept-level graph models [27]. Early MOOC course-recommendation systems such as that of Jing and Tang [2] already integrated multiple learner signals to this end.

Table 1. Comparison of recommendation approaches for MOOC course recommendation.

Family	Key idea	Strengths	Limitations	Example works
Collaborative filtering	Latent factors from co-interaction	Domain-agnostic, accurate with data	Sparsity, cold-start	[5], [7], [11], [12]
Content / knowledge-based	Match item features / domain rules	Eases cold-start, explainable	Over-specialisation	[13], [14], [15]
Sequential / session	Model order of interactions	Captures progression	Needs long histories	[9], [10], [16], [17]
Deep learning	Learned non-linear interactions	Flexible representations	Data-hungry, opaque	[18], [19], [20], [21]
Knowledge-graph-based	Reason over relational structure	Accurate, explainable, concept-level	KG construction cost	[23], [24], [25], [26], [27]
Reinforcement learning	Optimise long-term reward	Aligns with learning gain	Reward / sample efficiency	[22]
Hybrid	Combine multiple families	Offsets individual weaknesses	Design / tuning complexity	[2], [29]

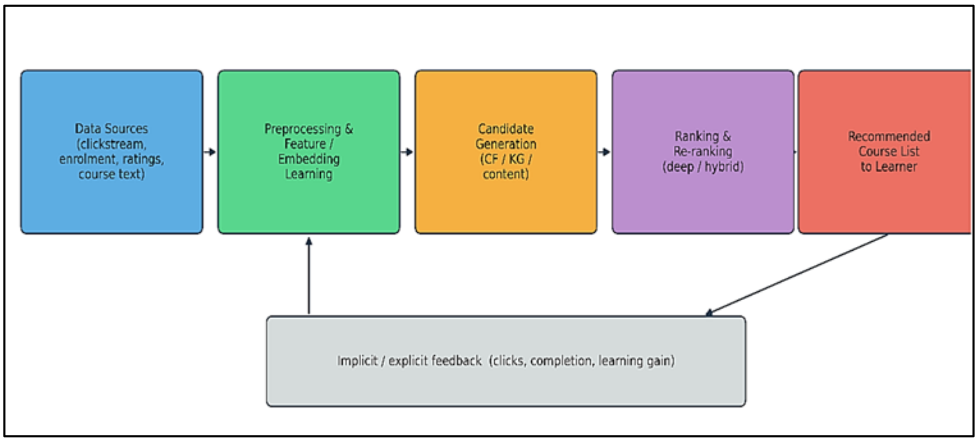

Fig. 2: Generic pipeline of a MOOC course-recommendation system with feedback loop.

IV. DATASETS AND PLATFORMS

Reproducible progress depends on shared benchmarks. The most widely used public resource is MOOCCube [1], a large repository released to support MOOC-oriented research. It links courses, concepts, videos and anonymised learner behaviour drawn from XuetangX. Its successor MOOCCubeX [30] expands the resource into a knowledge-centred repository designed for adaptive learning, with millions of behavioural records and fine-grained concept annotations that enable concept recommendation and knowledge tracing. XuetangX itself, one of China's largest MOOC platforms, is the source for much of this data and for several dropout and recommendation studies [4], [22]. Beyond the XuetangX ecosystem, research also draws on interaction logs and course catalogues from Coursera and edX, and KDD-Cup-style MOOC challenges have periodically released dropout and behaviour datasets. Table 2 summarises representative resources. One recurring caveat is that many platform datasets are proprietary or access-restricted, which limits cross-study comparability and keeps open repositories important.

Table 2. Representative datasets and platforms used in MOOC recommendation research.

Dataset / resource	Platform	Scale / content	Typical use
MOOCCube [1]	XuetangX	Courses, concepts, videos, learner logs	Course / concept rec.
MOOCCubeX [30]	XuetangX	Knowledge-centred, millions of records	Adaptive learning, KT
XuetangX logs [4], [22]	XuetangX	Large-scale clickstream / enrolment	Dropout, course rec.
Coursera / edX data	Coursera, edX	Catalogues, ratings, interaction logs	CF / content rec.
KDD-Cup MOOC sets	XuetangX et al.	Behaviour for dropout challenge	Dropout prediction

V. EVALUATION METRICS

Course recommendation is evaluated predominantly with top-K ranking metrics. Precision@K and Recall@K measure, respectively, the fraction of recommended items that are relevant and the fraction of relevant items recommended; F1@K is their harmonic mean. Hit-Rate@K records whether at least one relevant item appears in the top K. Rank-sensitive measures are preferred when position matters. Normalized Discounted Cumulative Gain (NDCG) discounts relevance by rank, Mean Average Precision (MAP) averages precision across recall levels, and Mean Reciprocal Rank (MRR) rewards placing the first relevant item near the top. These measures derive from classical information-retrieval and recommendation evaluation [6]. They dominate reported results across the families surveyed above.

Accuracy-oriented metrics capture only part of the educational objective. A growing line of work argues for learning-centric evaluation: course completion, certification, persistence and measured learning gain, as well as beyond-accuracy properties such as diversity, novelty and coverage. A recommendation that maximises immediate clicks may not maximise learning. Since dropout is the dominant failure mode in MOOCs [3], [4], aligning offline ranking metrics with downstream learning outcomes remains an open methodological problem. Table 3 summarises the principal metric groups.

Table 3. Evaluation metrics used in MOOC course-recommendation research.

Metric	Group	What it captures
Precision@K / Recall@K / F1@K	Accuracy	Relevance of the top-K list
Hit-Rate@K	Accuracy	Any relevant item in top K
NDCG@K	Ranking	Rank-discounted relevance
MAP	Ranking	Precision averaged over recall
MRR	Ranking	Position of first relevant item
Diversity / novelty / coverage	Beyond-accuracy	Breadth and freshness of list
Completion / learning gain	Learning-centric	Downstream educational outcome

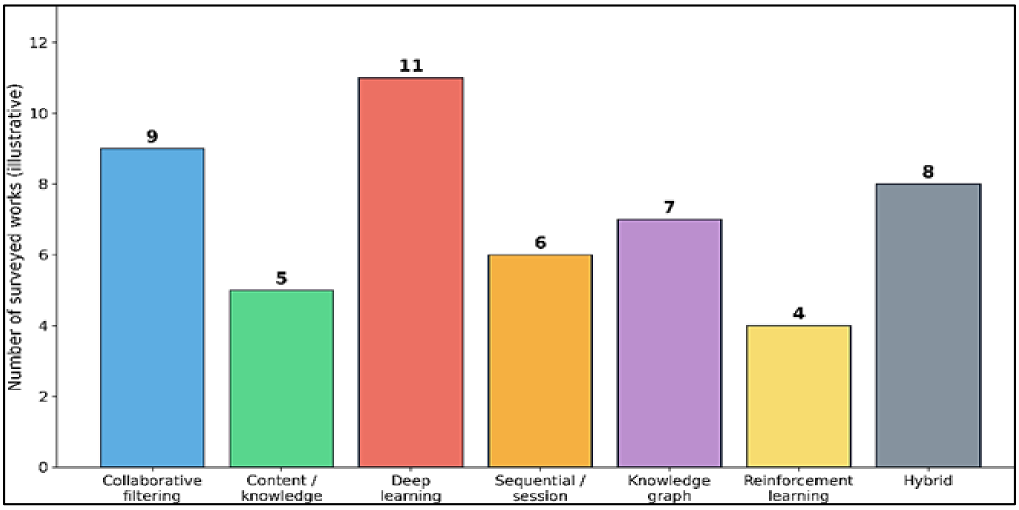

Fig. 3: Illustrative distribution of methodological families across the surveyed literature.

VI. OPEN CHALLENGES

Cold-start and data sparsity persist as the most cited obstacles. Most learners touch only a few of many thousands of courses, so the interaction matrix is almost entirely empty, and new learners or courses arrive without history [8]. Content, knowledge-graph and cross-domain signals are the common remedies, though none fully resolves the problem. Scalability is a related concern. Platforms serve millions of learners, so candidate generation and ranking must stay efficient, which has favoured precomputable item similarities [7] and lightweight graph models [25].

Integrating dropout prediction with recommendation is an increasingly prominent challenge. A poorly matched course raises the risk of disengagement, so jointly modelling who is likely to leave and what to recommend offers a path to retention-aware systems [3], [4]. Explainability and trust matter acutely in education, where learners and instructors benefit from understanding why a course was suggested; knowledge-graph models give a structural basis for such explanations [26], [27]. Fairness and diversity raise a further risk. Popularity-driven recommenders can entrench advantage, under-serve niche subjects or disadvantage particular learner groups, which motivates explicit diversity and fairness objectives.

The deepest challenge is pedagogical alignment. General recommenders optimise relevance, but education requires respecting the learner's evolving knowledge state and recommending material within the zone of proximal development, neither trivially easy nor unreachably hard [31]. This connects recommendation to knowledge tracing, the modelling of learner mastery over time, from the classical Bayesian formulation of Corbett and Anderson [32] to Deep Knowledge Tracing [33]. A recommender that ignores mastery may suggest courses for which a learner is unprepared. Privacy is a final, growing constraint: behavioural learning data are sensitive, and centralised collection conflicts with data-protection expectations, which motivates privacy-preserving designs.

VII. FUTURE DIRECTIONS

Several directions are gaining momentum, summarised on the evolutionary timeline in Fig. 4. The first is the use of large language models. Pretrained language models [34] and few-shot-capable large models [35] enable conversational, instruction-following course advisors. Such a system can interpret a learner's stated goals in natural language, reason over course descriptions, and justify its suggestions, moving recommendation from a ranked list toward an interactive dialogue. The second is knowledge-tracing-aware recommendation. An estimate of the learner's mastery from knowledge-tracing models [32], [33] conditions the recommender so that suggestions target the zone of proximal development [31] and advance a coherent learning path rather than maximising click-through alone.

A third direction is federated and privacy-preserving recommendation. It keeps sensitive behavioural data on learner devices or institutional servers and shares only model updates, aligning personalization with data-protection requirements. A fourth is the explicit treatment of fairness and diversity as first-class objectives, ensuring equitable exposure across subjects, providers and learner populations. Across all four, a unifying theme is the shift from accuracy-only optimisation toward learner-centred objectives that couple relevance with measured learning outcomes, and toward systems whose recommendations are transparent and accountable. Realising this agenda will require benchmarks that report learning-centric metrics alongside ranking accuracy, extending the repositories surveyed in Section IV.

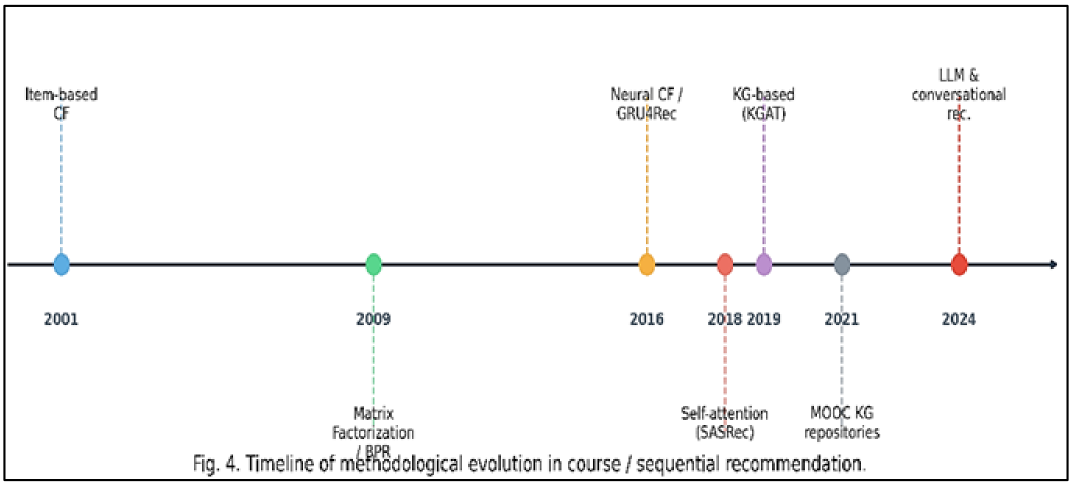

Fig. 4: Timeline of methodological evolution in course / sequential recommendation.

VIII. CONCLUSION

MOOC platforms have made education abundantly available, but abundance without guidance leaves learners overloaded and prone to dropout. This review traced the development of course-recommendation systems for MOOCs. The path runs from memory-based collaborative filtering and matrix factorization, through content, knowledge-based, sequential, deep-learning, knowledge-graph and reinforcement-learning methods, to hybrid and concept-level designs. The survey covered the datasets and platforms, principally the MOOCCube family and XuetangX, that enable reproducible study, and the ranking and learning-centric metrics by which systems are judged. The synthesis of open challenges, cold-start, sparsity, dropout integration, scalability, explainability, fairness, pedagogical alignment and privacy, indicates a field maturing from a purely accuracy-driven discipline toward a learner-centred one. The most promising directions, large-language-model and conversational recommenders, knowledge-tracing-aware systems, federated and privacy-preserving recommendation, and fairness, point toward recommenders that not only predict what a learner will click but help a learner genuinely progress. This consolidated reference is offered in support of that goal.

REFERENCES

[1] J. Yu, G. Luo, T. Xiao, Q. Zhong, Y. Wang, W. Feng, J. Luo, C. Wang, L. Hou, J. Li, Z. Liu, and J. Tang, "MOOCCube: A large-scale data repository for NLP applications in MOOCs," in *Proc. 58th Annu. Meeting Assoc. Comput. Linguistics (ACL)*, 2020, pp. 3135–3142.

[2] X. Jing and J. Tang, "Guess you like: Course recommendation in MOOCs," in *Proc. Int. Conf. Web Intelligence (WI)*, 2017, pp. 783–789.

[3] M. Kloft, F. Stiehler, Z. Zheng, and N. Pinkwart, "Predicting MOOC dropout over weeks using machine learning methods," in *Proc. EMNLP Workshop on Modeling Large Scale Social Interaction in Massively Open Online Courses*, 2014, pp. 60–65.

[4] W. Feng, J. Tang, and T. X. Liu, "Understanding dropouts in MOOCs," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, pp. 517–524, 2019.

[5] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, vol. 42, no. 8, pp. 30–37, 2009.

[6] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Trans. Knowl. Data Eng.*, vol. 17, no. 6, pp. 734–749, 2005.

[7] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, "Item-based collaborative filtering recommendation algorithms," in *Proc. 10th Int. Conf. World Wide Web (WWW)*, 2001, pp. 285–295.

[8] I. Schein, A. Popescul, L. H. Ungar, and D. M. Pennock, "Methods and metrics for cold-start recommendations," in *Proc. 25th ACM SIGIR Conf. Res. Develop. Inf. Retrieval (SIGIR)*, 2002, pp. 253–260.

[9] Hidasi, A. Karatzoglou, L. Baltrunas, and D. Tikk, "Session-based recommendations with recurrent neural networks," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2016.

[10] S. Wu, Y. Tang, Y. Zhu, L. Wang, X. Xie, and T. Tan, "Session-based recommendation with graph neural networks," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, pp. 346–353, 2019.

[11] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme, "BPR: Bayesian personalized ranking from implicit feedback," in *Proc. 25th Conf. Uncertainty Artif. Intell. (UAI)*, 2009, pp. 452–461.

[12] S. Rendle, "Factorization machines," in *Proc. IEEE Int. Conf. Data Mining (ICDM)*, 2010, pp. 995–1000.

[13] P. Lops, M. de Gemmis, and G. Semeraro, "Content-based recommender systems: State of the art and trends," in *Recommender Systems Handbook*. Boston, MA, USA: Springer, 2011, pp. 73–105.

[14] M. J. Pazzani and D. Billsus, "Content-based recommendation systems," in *The Adaptive Web*. Berlin, Germany: Springer, 2007, pp. 325–341.

[15] L. Pan, C. Li, J. Li, and J. Tang, "Prerequisite relation learning for concepts in MOOCs," in *Proc. 55th Annu. Meeting Assoc. Comput. Linguistics (ACL)*, 2017, pp. 1447–1456.

[16] W.-C. Kang and J. McAuley, "Self-attentive sequential recommendation," in *Proc. IEEE Int. Conf. Data Mining (ICDM)*, 2018, pp. 197–206.

[17] F. Sun, J. Liu, J. Wu, C. Pei, X. Lin, W. Ou, and P. Jiang, "BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer," in *Proc. 28th ACM Int. Conf. Inf. Knowl. Manage. (CIKM)*, 2019, pp. 1441–1450.

[18] S. Zhang, L. Yao, A. Sun, and Y. Tay, "Deep learning based recommender system: A survey and new perspectives," *ACM Comput. Surv.*, vol. 52, no. 1, pp. 1–38, 2019.

[19] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, "Neural collaborative filtering," in *Proc. 26th Int. Conf. World Wide Web (WWW)*, 2017, pp. 173–182.

[20] P. Covington, J. Adams, and E. Sargin, "Deep neural networks for YouTube recommendations," in *Proc. 10th ACM Conf. Recommender Syst. (RecSys)*, 2016, pp. 191–198.

[21] Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention Is All You Need," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2017, pp. 5998–6008.

[22] J. Zhang, B. Hao, B. Chen, C. Li, H. Chen, and J. Sun, "Hierarchical reinforcement learning for course recommendation in MOOCs," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, pp. 435–442, 2019.

[23] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2017.

[24] X. Wang, X. He, M. Wang, F. Feng, and T.-S. Chua, "Neural graph collaborative filtering," in *Proc. 42nd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval (SIGIR)*, 2019, pp. 165–174.
[25] X. He, K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang, "LightGCN: Simplifying and powering graph convolution network for recommendation," in *Proc. 43rd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval (SIGIR)*, 2020, pp. 639–648.
[26] X. Wang, X. He, Y. Cao, M. Liu, and T.-S. Chua, "KGAT: Knowledge graph attention network for recommendation," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, 2019, pp. 950–958.
[27] J. Gong, S. Wang, J. Wang, W. Feng, H. Peng, J. Tang, and P. S. Yu, "Attentional graph convolutional networks for knowledge concept recommendation in MOOCs in a heterogeneous view," in *Proc. 43rd Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval (SIGIR)*, 2020, pp. 79–88.
[28] S. Wu, F. Sun, W. Zhang, X. Xie, and B. Cui, "Graph neural networks in recommender systems: A survey," *ACM Comput. Surv.*, vol. 55, no. 5, pp. 1–37, 2023.
[29] R. Burke, "Hybrid recommender systems: Survey and experiments," *User Model. User-Adapt. Interact.*, vol. 12, no. 4, pp. 331–370, 2002.
[30] J. Yu, Y. Wang, Q. Zhong, G. Luo, Y. Mao, K. Sun, W. Feng, W. Xu, S. Cao, K. Zeng, Z. Yao, L. Hou, Y. Lin, P. Li, J. Zhou, B. Xu, J. Li, J. Tang, and M. Sun, "MOOCCubeX: A large knowledge-centered repository for adaptive learning in MOOCs," in *Proc. 30th ACM Int. Conf. Inf. Knowl. Manage. (CIKM)*, 2021, pp. 4643–4652.
[31] L. S. Vygotsky, *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA, USA: Harvard Univ. Press, 1978.
[32] T. Corbett and J. R. Anderson, "Knowledge tracing: Modeling the acquisition of procedural knowledge," *User Model. User-Adapt. Interact.*, vol. 4, no. 4, pp. 253–278, 1994.
[33] Piech, J. Bassen, J. Huang, S. Ganguli, M. Sahami, L. J. Guibas, and J. Sohl-Dickstein, "Deep knowledge tracing," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2015, pp. 505–513.
[34] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Human Lang. Technol. (NAACL-HLT)*, 2019, pp. 4171–4186.
[35] T. B. Brown, B. Mann, N. Ryder, *et al*., "Language models are few-shot learners," in *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 33, 2020, pp. 1877–1901.