PEER-REVIEWED JOURNAL OF COMPUTER SCIENCE (PRJCS)

ISSN Online: 3139-5082 Open Access Scholarly Journal

Ontology-Based Recommender Systems for Online Learning Platforms: A Review

Saritha E[1], Dr. B Kalpana[2]

[1] Research Scholar, Department of Computer Science, Avinashilingam Institute for Home Science and Higher Education for Women, Coimbatore, India.
[2] Professor & HOD, Department of Computer Science, Avinashilingam Institute for Home Science and Higher Education for Women, Coimbatore, India.

Article information

Received: 30th May 2026
Received in revised form: 18th June 2026
Accepted: 27th June 2026
Available online: 9th July 2026

Volume: 1
Issue: 7
DOI: https://doi.org/10.5281/zenodo.21232379

Abstract

The rapid growth of massive open online courses, learning management systems, and open educational repositories has produced an abundance of learning resources that overwhelms learners and complicates the selection of suitable content. Recommender systems have become a standard response to this information overload, yet the classical paradigms of collaborative filtering, content-based filtering, knowledge-based filtering, and their hybrids struggle with cold start, data sparsity, and a persistent semantic gap between raw interaction data and pedagogical meaning. This paper reviews how ontologies, formal and shareable specifications of a domain expressed in languages such as RDF and OWL, have been used to address these limitations in educational recommenders. It surveys the background of recommender paradigms and their weaknesses, explains what an ontology is and how domain, learner, and pedagogical ontologies encode knowledge, and analyses the mechanisms through which ontologies support semantic similarity, logical reasoning, learner modelling, prerequisite and zone of proximal development reasoning, and learning path sequencing. A taxonomy of ontology-based approaches is proposed and representative systems are compared, followed by a discussion of evaluation methods and datasets. The review then examines open challenges, including ontology construction cost, scalability, dynamic knowledge, cold start, evaluation validity, and explainability. It closes by outlining future directions that combine large language models with knowledge graphs, hybrid neuro-symbolic reasoning, and federated privacy-preserving personalization. Throughout, the aim is an honest synthesis of what prior work reports rather than any new empirical claim.

Keywords: - Ontology, Recommender Systems, E-Learning, Semantic Web, Knowledge Graph, Learner Modelling, Personalization, Learning Paths.

.

I. INTRODUCTION

Online learning platforms have expanded at a remarkable pace. Massive open online courses (MOOCs), institutional learning management systems (LMS), and open educational resource repositories now host quantities of material that far exceed what any single learner can inspect. This abundance produces information overload, a situation in which the effort required to locate relevant content becomes a barrier to learning rather than an aid. Recommender systems emerged in commercial settings as a practical remedy for exactly this problem, filtering large catalogues down to a personalized shortlist of items likely to interest a given user [1]. Their migration into education is natural, because learners face the same difficulty of choice while carrying the additional burden that a poor selection can waste study time and damage motivation.

The general theory and practice of recommendation are well documented, and comprehensive treatments describe the main families of algorithms together with their design trade-offs [2]. In the technology enhanced learning community, dedicated analyses have shown that recommenders are increasingly deployed to support learners, teachers, and course designers, but that the educational setting imposes constraints not present in retail or entertainment [3], [4]. A recommendation in a learning context should respect prerequisite structure, competency goals, and the learner current level of understanding, not merely predict what the learner might click. This pedagogical dimension motivates a large body of work that enriches recommenders with explicit domain knowledge.

Ontologies are a principal vehicle for injecting such knowledge. By providing a formal, shareable description of concepts and their relationships, an ontology allows a recommender to reason about the meaning of resources and learners rather than treating them as opaque identifiers. Several focused reviews have surveyed ontology-based recommenders for e-learning and have argued that they mitigate cold start and improve semantic relevance, while raising concerns about construction cost and maintenance [5], [6]. This paper synthesizes that literature. It reviews recommender paradigms and their limits, explains the ontology foundations, analyses the mechanisms and a taxonomy of ontology-based educational recommenders, compares representative systems and evaluation practice, and discusses open challenges and future directions. The account is descriptive and reports what prior work claims, without presenting new experiments.

Scoping is important because recommendation in education overlaps with adjacent tasks such as adaptive tutoring, curriculum planning, and search, and because the term ontology is sometimes used loosely for any structured metadata. This review restricts attention to systems that make personalized suggestions of learning resources, activities, peers, or paths, and that use an ontology in the strict sense of a formal conceptualization with defined classes and relations. Within that scope it draws primarily on dedicated surveys of ontology-based e-learning recommenders and on the foundational recommender and Semantic Web literature they build upon [2], [5], [6]. The remaining sections proceed from paradigms and their limits, to ontology foundations and mechanisms, to a taxonomy, representative systems, evaluation, challenges, and future work.

II. RECOMMENDER SYSTEMS: PARADIGMS AND LIMITATIONS

A. Collaborative Filtering

Collaborative filtering (CF) predicts a learner interest in an item from the behaviour of similar users or the co-occurrence patterns of similar items. Neighbourhood methods, and in particular item-based schemes, compute similarities from a user by item rating matrix and recommend items close to those a learner has already valued [7]. Model-based CF, most influentially matrix factorization, projects users and items into a shared latent space so that a dot product estimates preference, an approach that rose to prominence during large public recommendation competitions [8]. CF is attractive because it needs no description of item content and can surface unexpected but relevant material. Its weaknesses are equally well known. It suffers from data sparsity, because most learners interact with only a small fraction of the catalogue, and from the cold start problem, because a new learner or a newly added resource has no history from which to compute similarity.

B. Content-Based and Knowledge-Based Filtering

Content-based filtering recommends items whose descriptions resemble those a learner has previously preferred, using features drawn from text, metadata, or tags [9]. It handles new items gracefully, since a fresh resource can be described immediately, and it offers a degree of transparency because recommendations can be traced to shared features. However, content-based methods tend to overspecialize, recommending more of the same, and they depend on the quality of feature extraction, which for educational material rarely captures difficulty, prerequisite dependencies, or learning objectives. Knowledge-based recommenders instead reason over explicit rules and constraints that relate user requirements to item properties, which makes them robust to cold start but dependent on a knowledge base that must be authored and maintained.

C. Hybrid and Deep Approaches

Because each paradigm has complementary strengths, hybrid recommenders combine them, for example by weighting, switching, cascading, or feature augmentation, so that the components compensate for one another [10]. More recently, deep learning has been applied across all paradigms to model nonlinear interactions and rich content, and comprehensive surveys catalogue the resulting architectures and their perspectives [11]. Yet neither hybridization nor deep representation learning removes the underlying semantic gap. Interaction logs and shallow features do not encode why a resource is appropriate for a learner at a particular moment, whether a prerequisite has been met, or how two topics relate conceptually. This gap between statistical signal and pedagogical meaning is precisely what ontology-based methods set out to close, and it recurs throughout the reviews of educational recommenders [5], [6].

Two limitations deserve emphasis because ontologies address them directly. The first is cold start, which appears in three forms: a new learner with no history, a new resource with no ratings, and a new system with no accumulated data at all. Behaviour-driven methods are weakest exactly when a platform launches a course or admits a new cohort, which is when good guidance matters most. The second is the semantic gap. Two resources may cover the same concept with different words, or a beginner and an expert may click the same item for opposite reasons, and neither situation is visible to a model that sees only identifiers and counts. A recommender that lacks a representation of concepts, prerequisites, and objectives cannot tell that one resource must precede another, or that a highly rated advanced tutorial is unsuitable for a novice. These are not incidental flaws but structural consequences of learning from interaction alone, and they explain why the field turned to explicit knowledge. Table 1 summarizes the paradigms and their principal limitations.

Table 1. Comparison of classical Recommender Paradigms and Their Limitations in Education

Paradigm	**Core idea**	**Key strength**	**Main limitation**
Collaborative filtering	Recommend from behaviour of similar users or items	No content description needed, finds novel items	Sparsity and cold start for new users and items
Content-based	Match item descriptions to learner profile	Handles new items, offers feature-level transparency	Overspecialization, weak on difficulty and prerequisites
Knowledge-based	Reason over rules and constraints on requirements	Robust to cold start, encodes domain logic	Costly knowledge acquisition and maintenance
Hybrid	Combine two or more paradigms	Compensates for individual weaknesses	Added complexity, semantic gap can remain

III. ONTOLOGIES AND THE SEMANTIC WEB

An ontology, in the sense used across knowledge engineering and the Semantic Web, is a formal, explicit specification of a shared conceptualization [12]. It defines the classes, relations, functions, and constraints that describe a domain in a way that both people and machines can interpret consistently. This modelling view of knowledge, in which reusable structures capture the vocabulary of a field, grew out of the wider discipline of knowledge engineering and its methods for building knowledge-based systems [13]. The value of an ontology lies in shared understanding: once concepts and their relationships are agreed and formalized, different systems can exchange data and reasoning without ambiguity, and inference engines can derive facts that were not stated explicitly.

The Semantic Web provided the standards that made ontologies practical at scale, envisioning a web of data whose meaning is machine-processable [14]. Its stack includes the Resource Description Framework (RDF) for expressing statements as subject, predicate, object triples, RDF Schema for lightweight class and property hierarchies, and the Web Ontology Language (OWL) for richer axioms such as class equivalence, disjointness, and property characteristics that support automated reasoning. SPARQL provides a query language over RDF graphs. Together these give educational systems a common substrate in which learning content, learner attributes, and pedagogical rules can be represented as linked data and queried or reasoned over uniformly.

In the educational setting, three kinds of ontology recur. A domain ontology models the subject matter to be learned, for example the concepts of a programming language or a biology curriculum and the relationships, such as is a part of and is a prerequisite of, that connect them. A learner ontology models the student, including background, goals, preferences, prior knowledge, and progress. A pedagogical ontology models instructional constructs such as learning objectives, competencies, activity types, and assessment. Standardized metadata schemes complement these ontologies, most prominently the IEEE Learning Object Metadata standard, which specifies a structured vocabulary for describing learning objects and is widely used to annotate educational resources [15]. Competency ontologies extend the picture by describing skills and their dependencies, allowing a recommender to align resources with the abilities a learner still needs to acquire. Reviews of the field emphasize that this layered ontological modelling is what allows educational recommenders to reason pedagogically rather than only statistically [5], [6].

A practical design tension runs through this modelling. Richer axioms give a reasoner more to work with, but expressive OWL fragments can make inference slow or, in the limit, intractable, which is why the standard defines profiles that trade expressiveness for guaranteed efficient reasoning. Educational systems therefore choose a level of formality suited to their needs, from lightweight taxonomies that mainly support similarity, to rule-augmented ontologies that encode prerequisite and remediation policy. The same tension shapes ontology reuse: adopting an established vocabulary such as a standard metadata scheme lowers construction cost and improves interoperability, while a bespoke ontology fits a specific curriculum more tightly but is harder to share. These choices are not merely technical, since they determine how much genuinely pedagogical reasoning a recommender can perform and how portable its knowledge will be across courses and institutions.

IV. MECHANISMS: HOW ONTOLOGIES SUPPORT EDUCATIONAL RECOMMENDATION

Ontologies enter educational recommenders through several complementary mechanisms. The first is semantic similarity. Instead of comparing items by shared keywords, a system can measure how close two concepts are within the ontology graph, using path length, shared ancestors, or information content. This lets a recommender relate resources that use different terminology but describe the same idea, and it enriches both learner and item profiles with concepts inferred from the taxonomy. Hybrid learning material recommenders have combined such semantic profiles with collaborative signals and with the sequential patterns of previously accessed material to produce ordered suggestions [16], [17].

The second mechanism is reasoning. Because OWL ontologies carry logical axioms, a reasoner can infer implicit facts, check consistency, and apply rules. In a recommender, rules can express pedagogical policy, for example that an advanced module should be withheld until its prerequisites are satisfied, or that a learner who failed an assessment should be offered remedial content. This rule-based filtering narrows candidate items to those that are pedagogically admissible before ranking them. The third mechanism is learner modelling. An ontology gives the learner profile structure and meaning, so that observed behaviour updates a semantically rich model of knowledge state, goals, and preferences rather than a flat vector, which in turn supports more targeted recommendation and clearer explanations.

The fourth mechanism concerns sequencing and learning paths. Educational recommendation is often not about a single item but about the next best step in a trajectory toward a goal. Prerequisite relations in a domain ontology define a partial order over content, and reasoning over this order lets a system assemble a coherent path, avoid gaps, and adapt when a learner struggles. This connects to the pedagogical notion of the zone of proximal development, the band of tasks a learner can accomplish with appropriate support, since prerequisite and difficulty annotations let a recommender aim at material that is challenging but reachable. Ontology-based systems have used exactly this structure to generate personalized sequences and to recommend the next resource in a course [18], [19]. Reviews note that prerequisite and competency reasoning is among the clearest advantages ontologies bring over purely data-driven recommenders [5].

V. A TAXONOMY OF ONTOLOGY-BASED EDUCATIONAL RECOMMENDERS

Surveying the literature suggests a taxonomy along the dominant role the ontology plays. In semantic enrichment approaches, the ontology mainly supplies similarity and profile expansion that feed an otherwise conventional content-based or collaborative engine. In reasoning and rule-based approaches, an inference layer filters or ranks candidates according to logical constraints and pedagogical policy. In learner-modelling approaches, the ontology structures the student model and drives adaptation. In learning-path approaches, prerequisite and competency relations generate ordered sequences. Finally, hybrid ontology plus machine learning approaches combine ontological knowledge with statistical methods such as collaborative filtering, clustering, or sequential pattern mining, and improved hybrid designs of this kind have reported gains on online learning resource recommendation tasks [20]. These categories are not mutually exclusive, and many systems occupy more than one, but the taxonomy clarifies where the ontology contributes.

Figure 1 sketches the representative growth of publications on ontology-based educational recommenders over the last decade and a half, reflecting the sustained interest documented by the surveys rather than a precise bibliometric count. Table 2 details the taxonomy, associating each category with the ontology role, a typical technique, and representative lines of work identified in the reviewed literature.

Fig. 1. Representative (schematic) growth in publications on ontology-based educational recommenders. Values are illustrative of the upward trend reported across the surveyed reviews, not an exact bibliometric measurement.

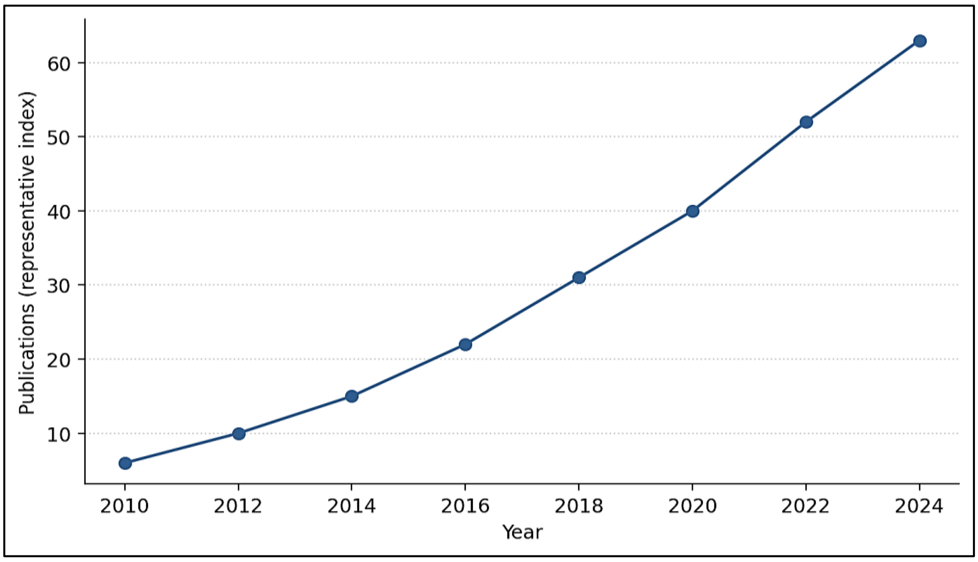

Table 2. Taxonomy of Ontology-Based Approaches in Educational Recommendation

Category	Role of the ontology	Typical technique	Representative focus in the literature
Semantic enrichment	Similarity and profile expansion	Concept similarity, feature augmentation	Enriched content and collaborative profiles
Reasoning and rules	Inference and constraint filtering	OWL reasoning, rule-based filtering	Pedagogically admissible candidate sets
Learner modelling	Structured student model	Ontological user profiles, clustering	Adaptation to knowledge state and style

Learning paths	Prerequisite and competency ordering	Sequencing, sequential pattern mining	Next-resource and path recommendation
Hybrid ontology plus ML	Knowledge prior for learning models	CF, clustering, pattern mining with ontology	Improved accuracy and cold-start handling

VI. REPRESENTATIVE SYSTEMS FROM THE LITERATURE

Several concrete systems illustrate the taxonomy. Semantic recommendation frameworks for e-learning have combined a domain ontology with OWL rules so that learners are matched to materials that fit their field of interest and current requirements, using rule filtering as the recommendation mechanism [19]. Such rule and ontology based designs demonstrate how explicit pedagogical policy can be encoded and applied, and they are robust to cold start because they do not require prior ratings for a new learner.

Hybrid knowledge-based systems have coupled ontologies with sequential pattern mining and collaborative filtering. One influential design models the learner and learning resources in an ontology, uses collaborative filtering to predict interest, and mines sequential access patterns to order recommendations, reporting improved relevance over non-semantic baselines [18]. Related hybrids incorporate learner context and the sequence of previously accessed material, arguing that both the learner situation and the order of study should shape what is recommended next. Learner-preference and sequential-pattern models that operate over multidimensional descriptions of material similarly show how ontological structure and behavioural sequence can be fused [17]. Earlier personalization systems, though not always framed as ontological, established the pattern of identifying learning styles and mining behaviour to adapt recommendations, and they informed subsequent semantic designs [16].

More recent hybrid ontology-based approaches for online learning resource recommendation have refined these ideas, combining collaborative filtering with sequential pattern mining under an ontological representation and reporting measurable improvements on resource recommendation tasks [20]. Across these systems a common architecture appears: an ontology layer that represents domain, learner, and pedagogy, a reasoning or similarity layer that interprets it, and a recommendation layer that ranks or sequences candidates. The reviews consolidate these examples and note recurring benefits in cold-start behaviour and semantic relevance, alongside recurring costs in ontology construction and upkeep [5], [6].

It is worth reading these results with care. Reported improvements are usually relative to non-semantic baselines on the authors own data, and gains vary with the quality and coverage of the ontology, the size of the interaction log, and the metric used. Systems that fuse an ontology with collaborative filtering and sequential pattern mining tend to report the strongest results, which suggests that the ontology contributes most when it complements rather than replaces behavioural signal [18], [20]. At the same time, few studies isolate the effect of the ontology through controlled ablation, so it is not always clear how much of an improvement comes from the semantic layer as opposed to the hybrid machinery around it. An honest reading of the literature is that ontologies reliably help with cold start, prerequisite ordering, and explanation, that they can improve accuracy when well constructed, and that the magnitude of benefit is context dependent and not yet established on shared benchmarks.

VII. EVALUATION METHODS AND DATASETS

Evaluation of ontology-based educational recommenders remains one of the least standardized aspects of the field. Offline studies dominate, reporting accuracy or ranking metrics such as precision, recall, F measure, mean absolute error, and normalized discounted cumulative gain on historical interaction data. These measures are convenient and reproducible, but they capture agreement with past behaviour rather than learning benefit, and they can reward a system that recommends easy or already familiar content. User studies and expert evaluations complement offline metrics by assessing perceived usefulness, satisfaction, and pedagogical appropriateness, yet they are smaller in scale and harder to compare across papers. The technology enhanced learning literature has repeatedly called for evaluation that measures effects on learning itself, such as knowledge gain, completion, or time to competency, and has proposed frameworks and layered criteria for doing so [3], [4].

Datasets are similarly fragmented. Unlike the general recommender field, which shares large public benchmarks, educational recommenders often rely on proprietary logs from a single platform or course, on small curated corpora, or on synthetic scenarios, which limits comparability and reproducibility. Public MOOC and LMS interaction datasets exist and are increasingly used, but ontological annotations must frequently be added by the authors, and the domain, learner, and pedagogical ontologies differ from study to study. The reviews conclude that the absence of shared, ontology-annotated benchmarks is a significant obstacle to cumulative progress [5], [6]. Reproducibility compounds the difficulty, since papers rarely release the ontology, the annotated data, or the code, so an independent team cannot rerun an experiment or test a new method against a prior one on equal terms. A mature evaluation practice would combine three layers, namely offline accuracy for rapid iteration, controlled user studies for perceived quality and appropriateness, and longer field studies that measure learning outcomes over time, and it would report ontology coverage and construction effort so that reported gains can be weighed against their cost. Until such practice becomes common, comparative claims about ontology-based educational recommenders should be read as promising indications rather than settled conclusions. Figure 2 gives a representative picture of how evaluation practice has shifted, with a gradual increase in user-centred and learning-oriented evaluation alongside the continued prevalence of offline accuracy metrics.

Fig. 2. Representative (schematic) distribution of evaluation approaches in ontology-based educational recommender studies, contrasting earlier and more recent work. Percentages are illustrative of trends reported in the surveys, not a precise meta-analysis.

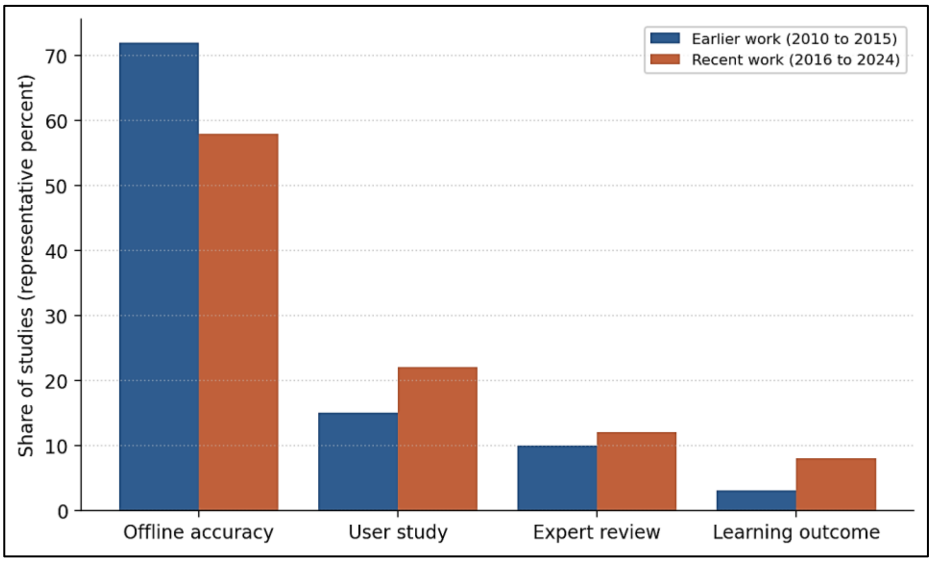

VIII. OPEN CHALLENGES

The first and most cited challenge is ontology construction cost. Building an accurate domain ontology, together with learner and pedagogical models, demands substantial expert effort, and the resulting artefacts must be validated. For fast moving subjects the cost recurs, because the ontology has to be revised as the field changes. This barrier explains why many systems remain research prototypes confined to a single course. The second challenge is scalability. Reasoning over large OWL ontologies can be computationally expensive, and coupling it with recommendation over many learners and resources raises questions of response time and infrastructure that the literature acknowledges but seldom resolves at platform scale.

The third challenge is dynamic knowledge. Curricula, resources, and learner goals change continuously, yet an ontology is a relatively static structure, so keeping it current is an ongoing maintenance problem. The fourth is cold start, which ontologies mitigate but do not eliminate. A knowledge-based layer can recommend for a new learner using stated goals and prerequisites, but the quality of those recommendations still depends on how completely the learner and domain are modelled. The fifth is evaluation validity, discussed above, where offline metrics may not reflect learning outcomes and shared benchmarks are scarce. The sixth is explainability. Ontologies are often promoted as a route to transparent, justifiable recommendations, and reasoning traces can indeed explain why an item was suggested, yet turning logical justifications into explanations that learners find clear and actionable is not automatic and remains underexplored. Knowledge graph based recommenders, which generalize the ontology idea to large interlinked entity graphs, face closely related challenges of construction, scale, and explanation, and surveys of that area document both the promise and the difficulty [21], [22].

Two further challenges are receiving growing attention. The first is trust and fairness. Because an ontology encodes human decisions about what concepts exist and how they relate, biases in its design can propagate into recommendations, for example by privileging certain topics, languages, or learner groups, and an authoritative-looking semantic justification can lend unwarranted confidence to a flawed suggestion. The second is integration and adoption. Many ontology-based recommenders are validated as standalone prototypes, yet real platforms already run their own catalogues, analytics, and identity systems, so embedding a reasoning layer without degrading responsiveness or privacy is a nontrivial engineering task. Together with construction cost, scalability, dynamic knowledge, cold start, evaluation validity, and explainability, these concerns describe why strong laboratory results have not yet translated into widespread deployment, and they set the agenda for the directions considered next.

IX. FUTURE DIRECTIONS

A prominent direction combines large language models (LLMs) with knowledge graphs and ontologies. LLMs offer broad language understanding and can help construct or populate ontologies, extract concepts and prerequisite relations from course text, and generate natural language explanations, while the ontology or knowledge graph supplies factual grounding, structure, and constraints that curb hallucination. Roadmaps for unifying LLMs and knowledge graphs describe complementary patterns, using graphs to enhance models and using models to build and complete graphs, that map directly onto educational recommendation needs [23]. In parallel, surveys of LLMs for recommendation catalogue discriminative and generative uses that could bring conversational, context-aware suggestion to learning platforms while retaining the semantic backbone an ontology provides [24]. Retrieval-augmented generation, in which a model conditions its output on documents or facts fetched at inference time, is a practical realization of this pattern for knowledge-intensive applications and offers a route to ground recommendation and explanation in a curated ontology or resource base [25].

A second direction is hybrid neuro-symbolic recommendation, in which neural representation learning and symbolic ontological reasoning operate together rather than in isolation. Such systems could learn preferences and content embeddings from data while enforcing prerequisite logic and pedagogical rules symbolically, combining the adaptivity of machine learning with the guarantees and transparency of explicit knowledge. This transparency connects to the wider agenda of explainable artificial intelligence in critical decision-making systems, where the ability to justify an automated choice is treated as a requirement rather than a convenience, a concern that applies squarely to recommendations that shape a learner path [26]. Knowledge graph based recommenders already point in this direction by embedding entities and relations while preserving graph structure, and the education-specific versions of these methods are a natural next step [22].

A third direction concerns privacy and decentralization. Learner data is sensitive, and centralizing detailed models across institutions raises legitimate concerns. Federated and privacy-preserving learning, in which models are trained across distributed data without moving the raw records, offers a way to personalize while protecting learners, and an ontology shared across sites can provide the common semantic frame that makes federated educational recommendation coherent. Taken together, these directions suggest a convergence in which ontologies and knowledge graphs supply meaning and pedagogy, neural and language models supply flexibility and scale, and privacy-preserving training supplies trust, addressing the very limitations that recur throughout the reviewed literature [5], [6], [21].

X. CONCLUSION

Ontology-based recommendation offers a principled answer to information overload on online learning platforms by supplying the semantic and pedagogical knowledge that purely statistical recommenders lack. This review traced the path from classical recommender paradigms and their cold-start and semantic-gap limits, through the ontology and Semantic Web foundations, to the mechanisms by which domain, learner, and pedagogical ontologies support semantic similarity, reasoning, learner modelling, and prerequisite-aware learning paths. A taxonomy organized the approaches, representative systems illustrated them, and a survey of evaluation practice exposed a reliance on offline accuracy and a shortage of shared, ontology-annotated benchmarks and learning-outcome measures.

The evidence from prior work is encouraging but qualified. Ontologies consistently help with cold start, semantic relevance, prerequisite reasoning, and explanation, yet they impose real costs in construction, maintenance, and scalability, and the field lacks the standardized evaluation that would let competing designs be compared fairly. The most promising way forward joins ontologies and knowledge graphs with large language models and neuro-symbolic methods, under privacy-preserving and federated training, so that meaning, adaptivity, and trust reinforce one another. If future systems pair explicit pedagogical knowledge with scalable learning and honest, learning-centred evaluation, ontology-based recommenders can move from research prototypes toward dependable components of everyday online education.

REFERENCES

[1] J. Bobadilla, F. Ortega, A. Hernando, and A. Gutierrez, “Recommender systems survey,” Knowledge-Based Systems, vol. 46, pp. 109–132, 2013.

[2] F. Ricci, L. Rokach, and B. Shapira, Eds., Recommender Systems Handbook, 2nd ed. New York, NY, USA: Springer, 2015.

[3] N. Manouselis, H. Drachsler, R. Vuorikari, H. Hummel, and R. Koper, “Recommender systems in technology enhanced learning,” in Recommender Systems Handbook, F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, Eds. Boston, MA, USA: Springer, 2011, pp. 387–415.

[4] H. Drachsler, K. Verbert, O. C. Santos, and N. Manouselis, “Panorama of recommender systems to support learning,” in Recommender Systems Handbook, 2nd ed., F. Ricci, L. Rokach, and B. Shapira, Eds. New York, NY, USA: Springer, 2015, pp. 421–451.

[5] J. K. Tarus, Z. Niu, and G. Mustafa, “Knowledge-based recommendation: A review of ontology-based recommender systems for e-learning,” Artificial Intelligence Review, vol. 50, no. 1, pp. 21–48, 2018.

[6] G. George and A. M. Lal, “Review of ontology-based recommender systems in e-learning,” Computers and Education, vol. 142, Art. no. 103642, 2019.

[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), Hong Kong, 2001, pp. 285–295.

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

[9] P. Lops, M. de Gemmis, and G. Semeraro, “Content-based recommender systems: State of the art and trends,” in Recommender Systems Handbook, F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, Eds. Boston, MA, USA: Springer, 2011, pp. 73–105.

[10] R. Burke, “Hybrid recommender systems: Survey and experiments,” User Modeling and User-Adapted Interaction, vol. 12, no. 4, pp. 331–370, 2002.

[11] S. Zhang, L. Yao, A. Sun, and Y. Tay, “Deep learning based recommender system: A survey and new perspectives,” ACM Computing Surveys, vol. 52, no. 1, pp. 1–38, 2019.

[12] T. R. Gruber, “A translation approach to portable ontology specifications,” Knowledge Acquisition, vol. 5, no. 2, pp. 199–220, 1993.

[13] R. Studer, V. R. Benjamins, and D. Fensel, “Knowledge engineering: Principles and methods,” Data and Knowledge Engineering, vol. 25, nos. 1–2, pp. 161–197, 1998.

[14] T. Berners-Lee, J. Hendler, and O. Lassila, “The Semantic Web,” Scientific American, vol. 284, no. 5, pp. 34–43, 2001.

[15] IEEE Standard for Learning Object Metadata, IEEE Std 1484.12.1-2002, IEEE, 2002.

[16] A. Klasnja-Milicevic, B. Vesin, M. Ivanovic, and Z. Budimac, “E-learning personalization based on hybrid recommendation strategy and learning style identification,” Computers and Education, vol. 56, no. 3, pp. 885–899, 2011.

[17] M. Salehi and I. N. Kamalabadi, “Hybrid recommendation approach for learning material based on sequential pattern of the accessed material and the learner’s preference tree,” Knowledge-Based Systems, vol. 48, pp. 57–69, 2013.

[18] J. K. Tarus, Z. Niu, and A. Yousif, “A hybrid knowledge-based recommender system for e-learning based on ontology and sequential pattern mining,” Future Generation Computer Systems, vol. 72, pp. 37–48, 2017.

[19] S. Shishehchi, S. Y. Banihashem, and N. A. M. Zin, “A proposed semantic recommendation system for e-learning: A rule and ontology based e-learning recommendation system,” in Proc. Int. Symp. Information Technology (ITSim), Kuala Lumpur, Malaysia, 2010, pp. 1–5.

[20] S. Shanshan, G. Mingjin, and L. Lijuan, “An improved hybrid ontology-based approach for online learning resource recommendations,” Educational Technology Research and Development, vol. 69, no. 5, pp. 2637–2661, 2021.

[21] A. Hogan et al., “Knowledge graphs,” ACM Computing Surveys, vol. 54, no. 4, pp. 1–37, 2021.

[22] Q. Guo et al., “A survey on knowledge graph-based recommender systems,” IEEE Transactions on Knowledge and Data Engineering, vol. 34, no. 8, pp. 3549–3568, 2022.

[23] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, and X. Wu, “Unifying large language models and knowledge graphs: A roadmap,” IEEE Transactions on Knowledge and Data Engineering, vol. 36, no. 7, pp. 3580–3599, 2024.

[24] L. Wu et al., “A survey on large language models for recommendation,” World Wide Web, vol. 27, no. 5, Art. no. 60, 2024.

[25] “Retrieval-augmented generation for knowledge-intensive language applications,” Peer-Reviewed Journal of Computer Science (PRJCS), E-ISSN 3139-5082, 2025.

[26] “Explainable artificial intelligence in critical decision-making systems,” Peer-Reviewed Journal of Computer Science (PRJCS), E-ISSN 3139-5082, 2025.