EURASIAN JOURNAL OF MATHEMATICAL AND COMPUTER APPLICATIONS
ISSN 2306–6172
Volume 14, Issue 1 (2026) 17 – 34

KNOWLEDGE AUGMENTED GENERATION FOR CURRICULUM PLANNING

Kozhanov M., Mako C., Poser V., Eigner G., Mosavi A.*

Abstract Interdisciplinary education is positioned as a strategic response to complex technological and societal challenges. However, outcome-driven synthesis of complete study plans remains difficult to scale under heterogeneous competency standards and strict academic regulations. Retrieval-Augmented Generation (RAG) improves factual grounding by retrieving external evidence, while Knowledge-Augmented Generation (KAG) extends RAG by integrating structured knowledge representations for relational reasoning and domain robustness. This paper introduces *Curriculum-KAG*, a new method and implemented information system that mirrors KAG onto the macro-level task of interdisciplinary curriculum synthesis. A curriculum knowledge base integrates (i) a vector index for semantic retrieval and (ii) a curriculum knowledge graph encoding prerequisites, domains, and regulatory constraints. Hybrid retrieval with graph expansion selects candidate courses and enforces prerequisite closure. Constrained synthesis is formulated as multi-objective optimization with strict verification, while bridge modules are generated only under an evidence constraint when critical learning outcomes remain weakly covered. A prototype is reported on a two-domain case study (IT + Forensics) for the "Cyber Investigator" programme, including an 8-semester plan (240 ECTS) and outcome coverage diagnostics (LO1–LO7, with LO5 at 65%). Measured evaluation against baselines indicates improvements in retrieval and mapping (Recall@20=0.91 ± 0.02, nDCG@20=0.88 ± 0.02, Macro-F1=0.85 ± 0.02) while preserving feasibility (0 violations) and reducing redundancy (0.41 ± 0.03).

Keywords: Curriculum generation, interdisciplinary programs, knowledge graphs, evidence-constrained synthesis, graph expansion, artificial intelligence, machine learning, deep learning, augmented learning, syllabus design, data science.

AMS Mathematics Subject Classification: 68T01, 68T05, 68T30.

DOI: 10.32523/2306-6172-2026-14-1-17-34

1 Introduction

Interdisciplinary education is recognized as a strategic response to complex societal and technological challenges that exceed the scope of single disciplines. Such curricula are structured to integrate heterogeneous knowledge domains, which foster critical thinking, analytical reasoning, and communicative competence [1–3]. Furthermore, interdisciplinary synthesis programs are empirically associated with enhanced research productivity and cross-domain collaboration, as coordinated funding initiatives are shown to stimulate integrative outputs [4]. Therefore, structured and outcome-driven curriculum design approaches are required in higher education contexts in which adaptability, innovation, and accountability are emphasized [5,6]. In addition, explicit alignment between competencies, course content, and assessment criteria is required to ensure academic rigor and transparency. Retrieval-Augmented Generation

[1]Corresponding Author.

(RAG) is defined as a generative approach in which external corpora are retrieved and incorporated into the generation process to improve factual grounding and contextual relevance [7]. Knowledge-Augmented Generation (KAG) extends RAG by integrating structured knowledge representations, such as knowledge graphs, to support relational reasoning and domain robustness [8,9]. In KAG architectures, vector-based retrieval is combined with graph-based semantic structures, which encode hierarchical and associative relations among concepts. Moreover, multi-source knowledge pruning is applied to remove redundant or irrelevant evidence, which enhances precision and reliability [10]. In educational contexts, KAG is aligned with psychometric calibration and difficulty control, which ensure that generated outputs correspond to intended cognitive levels and learning objectives [11, 12]. As a result, KAG is positioned as a method that supports evidence-based and cognitively aligned content production. Despite these technological advances, curriculum development in interdisciplinary programs is frequently conducted through manual expert deliberation, which is time-intensive and difficult to scale. Fragmentation of knowledge across departments is often observed, while coherent sequencing across domains is not systematically ensured [1]. Furthermore, competency-based standards differ across disciplines, which complicates alignment between learning outcomes, prerequisite structures, and assessment models [13, 14]. Consequently, the synthesis of outcome-driven interdisciplinary study plans remains methodologically under-specified. In other words, no formal correspondence between knowledge-augmented generation principles and curriculum construction procedures is explicitly articulated in the literature.

Although KAG is applied to micro-level educational tasks, such as question generation and educational question answering, its extension to macro-level curriculum synthesis remains limited. For example, KAQG integrates knowledge graphs with RAG to control question difficulty and cognitive complexity [15, 16]. Similarly, MDKAG employs a multimodal disciplinary knowledge graph to enhance evidence retrieval and answer relevance [9]. Knowledge-augmented instructional models are also validated in specific course contexts, such as TRIZ education, in which adaptive materials and feedback are generated [17]. However, a comprehensive method that transfers retrieval, graph reasoning, and constrained generation to the design of complete interdisciplinary study plans is not fully developed. Empirical evidence indicates that knowledge graph-driven educational agents facilitate interdisciplinary mapping of concepts, which supports cross-domain reasoning and instructional assistance [18]. In addition, interdisciplinary curriculum initiatives in product and media design education demonstrate that structured knowledge integration enhances competency development and innovation capacity [19]. STEM teacher education programs further report improvement of digital and scientific competencies when disciplinary boundaries are explicitly bridged [20].

Structural modeling techniques that combine Graph Neural Networks (GNN) with Long Short-Term Memory (LSTM) architectures are also shown to achieve measurable predictive performance for next-course recommendation, with Mean Reciprocal Rank (MRR) $= 0.2545$, Hits@10 $= 0.7222$, and next-course accuracy (Acc) $= 0.7561$ [21–23]. These results indicate that curriculum structures are computationally representable with acceptable validity. Nevertheless, a second research gap is identified in the absence of a multi-objective formulation that enforces structural validity constraints in automated curriculum generation. Existing KAG applications prioritize factual grounding and reasoning accuracy; however, regulatory constraints, prerequisite chains, credit distributions, and accreditation standards are not formally encoded in generative objectives. Comparative analyses of national and international curricula reveal inconsistencies in competency articulation and content emphasis, which further complicate automated synthesis [13, 24]. Therefore, evidence-centric retrieval alone is insufficient, as long as structural coherence and academic regulations are not explicitly constrained. A constrained synthesis approach that merges, reorders, and proposes bridge courses under formal academic rules remains insufficiently specified. In response to these limitations, Curriculum-KAG is introduced as a method that mirrors the KAG paradigm within interdisciplinary curriculum generation [25]. A curriculum knowledge base is constructed, which integrates a vector index for semantic retrieval with a curriculum knowledge graph that encodes disciplinary relations, prerequisite dependencies, competency mappings, and regulatory requirements. Hybrid retrieval with graph expansion is applied to ensure evidence-centric and relation-aware selection of candidate courses. Constrained plan synthesis is then performed, in which courses are merged, reordered, or newly proposed as bridge modules under explicit academic regulations. Novelty is established through a formal mapping between RAG and KAG components and curriculum generation operations, a multi-objective optimization formulation with structural validity constraints, and an evidence-based evaluation protocol with predefined acceptance thresholds. Validated modules from prior research are integrated, which includes document analysis through combined vectorization and machine learning, achieving reliable performance [26]. Therefore, we aim at implementing Curriculum-KAG to investigate the contribution of both the theoretical extension of knowledge-augmented generation to curriculum science and a reproducible method for outcome-driven interdisciplinary study plan synthesis under formal academic constraints.

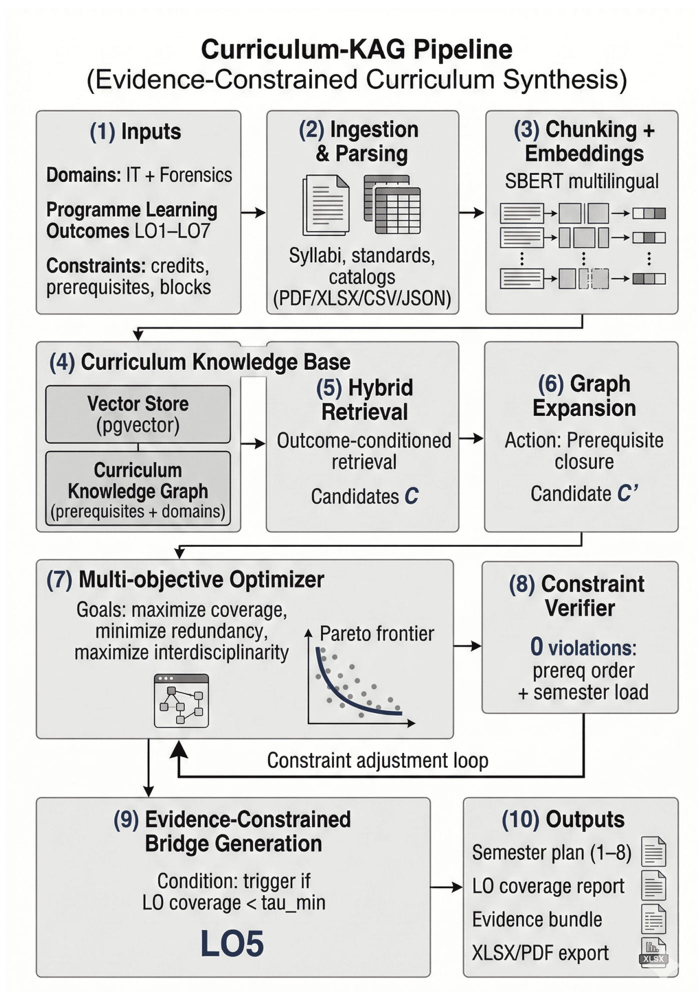

Figure 1: Curriculum-KAG pipeline, i.e., evidence-constrained curriculum synthesis.

Outcome-based education requires explicit alignment between intended learning outcomes, course content, and assessment. Constructive alignment and backward design motivate programme design from intended outcomes backward to learning activities and evaluation [27,28]. KAG architectures integrate vector retrieval with graph-based semantic structures for relational reasoning and domain robustness [8]. Knowledge graph-guided RAG explicitly uses KG structure for evidence retrieval and grounding [16]. GraphRAG constructs graph representations from corpora to enable local-to-global retrieval and summarization [10]. Curriculum analytics provides graph-based complexity measures for curricula and sequencing [24, 29].

2 Problem formulation

Let C denote the set of candidate courses and O denote the set of programme learning outcomes (PLOs). A directed prerequisite relation is represented by $E \subseteq C \times C$. A curriculum plan is represented by a binary selection vector $x \in \{0,1\}^{|C|}$ and a semester assignment $\mathrm{sem}(c) \in \{1, \dots, S\}$ for selected courses. A core artifact is the course–outcome coverage matrix $W \in [0,1]^{|C|\times|O|}$.

Hard constraints enforce credit targets and prerequisite order, including both minimum and maximum semester loads to ensure student status viability [30, 31]:

$$\sum_{c \in C} x_c \cdot \mathrm{cr}(c) = \mathrm{CR}_{\mathrm{target}}, \tag{1}$$

$$\forall (c_i, c_j) \in E: \quad x_{c_i} = x_{c_j} = 1 \Rightarrow \mathrm{sem}(c_i) < \mathrm{sem}(c_j), \tag{2}$$

$$\forall s: \quad \mathrm{CR}_{\min}(s) \le \sum_{c:\mathrm{sem}(c)=s} x_c \cdot \mathrm{cr}(c) \le \mathrm{CR}_{\max}(s). \tag{3}$$

Multi-objective synthesis (using a bounded cumulative sum for coverage to reflect knowledge accumulation [32, 33]):

$$\max\ f_{\mathrm{cov}}(x) = \frac{1}{|O|} \sum_{o \in O} \min\Big(1.0, \sum_{c:x_c=1} W_{c,o}\Big), \tag{4}$$

$$\min\ f_{\mathrm{red}}(x) = \frac{2}{|S_x|(|S_x|-1)} \sum_{c_i < c_j \in S_x} \mathrm{sim}(c_i, c_j), \quad S_x = \{c : x_c = 1\}, \tag{5}$$

$$\max\ f_{\mathrm{inter}}(x) = -\sum_{d \in \mathcal{D}} p_d(x) \log p_d(x), \tag{6}$$

where $\mathrm{sim}(\cdot,\cdot)$ is cosine similarity of course embeddings and $p_d(x)$ is the credit-weighted domain proportion.

3 Theoretical framework: AI-enabled curriculum design

In this section we consider the curriculum Knowledge Base, i.e., vector store plus curriculum KG. A Curriculum Knowledge Base (CKB) is constructed as $(\mathcal{D}, \mathcal{V}, \mathcal{G})$: $\mathcal{D}$ is a corpus of standards, syllabi, and catalogs; $\mathcal{V}$ is a chunk-level vector index; $\mathcal{G}$ is a curriculum knowledge graph. Texts are chunked and vectorized using multilingual Sentence-Transformers [34, 35]. The vector store is implemented in PostgreSQL with `pgvector`. The graph encodes prerequisites, domains, and regulation constraints, and is constructed via a semi-automated LLM-based relation extraction pipeline validated by domain experts to ensure scalability and accuracy [36].

(CKB) is designed as an integrated data structure that supports knowledge retrieval and reasoning for curriculum planning. The CKB is composed of three main components. A document corpus contains educational standards, course syllabi, and academic catalogues that serve as primary information sources. These documents are segmented into smaller text units and are transformed into numerical vector representations through multilingual Sentence Transformers. The vectors are stored in a vector database that is implemented in PostgreSQL with the pgvector extension, which allows efficient semantic similarity search across the corpus. In addition to the vector store, a curriculum knowledge graph is constructed in order to represent structured relations between courses and academic requirements. In this graph, prerequisite dependencies, disciplinary domains, and regulatory constraints are formally encoded. The graph structure is generated through a semi automated relation extraction process that is performed with large language models. The extracted relations are reviewed and validated by domain experts in order to ensure reliability and correctness. Through the integration of semantic retrieval from the vector index and relational structure from the knowledge graph, the CKB supports scalable and accurate information access for curriculum synthesis.

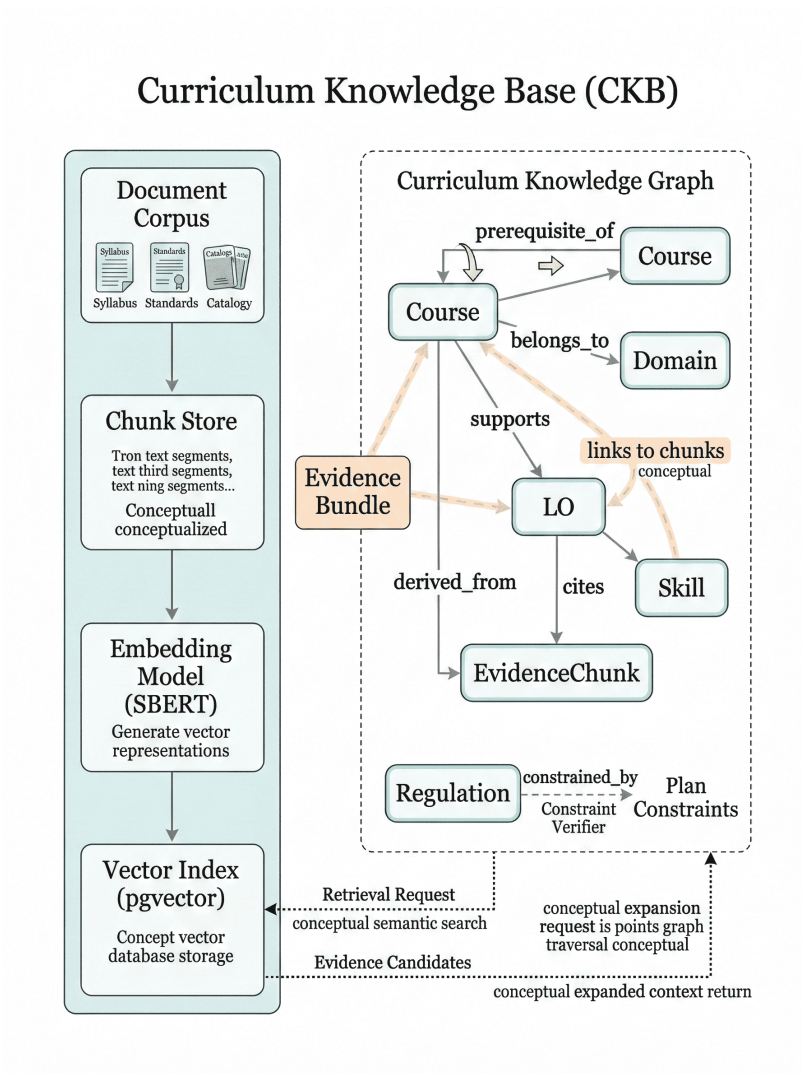

Figure 2: CKB: linkage between vector retrieval and curriculum KG reasoning, enabling evidence bundles and constraint-aware expansion.

4 Curriculum-KAG method: evidence-constrained curriculum synthesis

4.1 Hybrid retrieval with graph expansion

Let the curriculum knowledge graph be defined as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = C$ is the set of course nodes and $\mathcal{E}$ represents the directed prerequisite relations. We define the adjacency matrix $\mathbf{A} \in \{0,1\}^{|C|\times|C|}$, where $\mathbf{A}_{i,j} = 1$ if course c_i is a strict prerequisite for course c_j.

Given target outcomes O, the initial semantic retrieval returns a ranked candidate pool C', which can be represented as a binary indicator vector $\mathbf{v}^{(0)} \in \{0,1\}^{|C|}$. To ensure structural feasibility, graph expansion enforces prerequisite closure via the reachability matrix $\mathbf{R}$. The topological transitive closure of depth L is computed as [37]:

$$\mathbf{R} = \operatorname{sgn}\left(\sum_{k=1}^{L} \mathbf{A}^k\right), \tag{7}$$

where $\operatorname{sgn}(\cdot)$ is the element-wise signum function. The expanded candidate set C'', ensuring that all prerequisites of the selected courses are included, is obtained by mapping the initial vector:

$$\mathbf{v}^{(\text{closed})} = \operatorname{sgn}\left(\mathbf{R}^T \mathbf{v}^{(0)} \vee \mathbf{v}^{(0)}\right), \tag{8}$$

thus $C'' = \{c_i \in C \mid \mathbf{v}_i^{(\text{closed})} = 1\}$.

4.2 Coverage estimation and admissibility

Let $\phi : \mathcal{D} \to \mathbb{R}^d$ denote the dense embedding function mapping textual representations of a course $t(c)$ and an outcome $t(o)$ to a d-dimensional continuous space [34]. The semantic similarity is defined via the inner product of L_2-normalized vectors:

$$s_{\text{sem}}(c,o) = \frac{\langle \phi(t(c)), \phi(t(o)) \rangle}{\|\phi(t(c))\|_2 \|\phi(t(o))\|_2}. \tag{9}$$

To mitigate hallucination in dense retrieval, semantic signals are combined with exact-match lexical signals using the BM25 algorithm [38]. For a learning outcome represented as a set of query terms $Q_o = \{q_1, \ldots, q_n\}$, the lexical score against a course document c is:

$$s_{\text{lex}}(c,o) = \sum_{q_i \in Q_o} \text{IDF}(q_i) \cdot \frac{f(q_i,c)\cdot(k_1+1)}{f(q_i,c) + k_1 \cdot \left(1 - b + b \cdot \frac{|c|}{\text{avgdl}}\right)}, \tag{10}$$

where $f(q_i, c)$ is the term frequency, and k_1, b are standard hyper-parameters. The scores are Min-Max normalized to $\widetilde{s}_{\text{sem}}, \widetilde{s}_{\text{lex}} \in [0,1]$.

The hybrid coverage weight $W_{c,o}$ is subsequently computed as a convex combination ($\alpha = 0.75$):

$$W_{c,o} = \alpha \cdot \widetilde{s}_{\text{sem}}(c,o) + (1-\alpha)\cdot \widetilde{s}_{\text{lex}}(c,o). \tag{11}$$

Furthermore, evidence admissibility imposes a hard constraint derived from the knowledge graph:

$$\text{Adm}(c,o) \iff \exists e \in \mathcal{E} : \ \text{supports}(e,c,o), \tag{12}$$

ensuring that mapping is permitted only if structurally validated by the CKB.

4.3 Optimization, verifier, and bridge trigger

The multi-objective synthesis defined in Eq. (4)–(6) is solved using the Non-dominated Sorting Genetic Algorithm II (NSGA-II) [25]. Let $F(x) = [f_{\text{cov}}(x), -f_{\text{red}}(x), f_{\text{inter}}(x)]^T$ be the objective vector to be maximized. A curriculum plan $x^{(1)}$ strictly Pareto-dominates $x^{(2)}$ (denoted $x^{(1)} \succ x^{(2)}$) if and only if:

$$\forall m \in \{1,2,3\}: \; F_m(x^{(1)}) \geq F_m(x^{(2)}) \quad \wedge \quad \exists k \in \{1,2,3\}: \; F_k(x^{(1)}) > F_k(x^{(2)}). \tag{13}$$

The optimization (population size $N = 100$, maximum generations $G = 200$, crossover probability $p_c = 0.9$, mutation probability $p_m = 0.1$) maintains diversity via crowding distance and respects the hard constraints (Eq. 1–3) via penalty functions.

A deterministic verifier evaluates the Pareto-optimal plans. A bridge module generation is conditionally triggered if the minimum coverage of any critical learning outcome $o \in O^\star$ falls below a predefined threshold $\tau_{\min}$:

$$\min_{o \in O^\star} \left[\min\left(1.0, \sum_{c: x_c = 1} W_{c,o}\right) \right] < \tau_{\min}. \tag{14}$$

This ensures that interdisciplinary gaps are automatically detected and flagged for synthetic content augmentation.

5 Case study: “Cyber Investigator” (IT + Forensics)

5.1 Learning outcomes (LO1–LO7)

LO1. Analyze the architecture of information systems and networks to identify vulnerabilities and traces of cyber incidents.
LO2. Apply digital forensics methods to collect, preserve, and document digital evidence in compliance with procedural requirements.
LO3. Investigate cybercrimes using specialized forensic tools and analytical platforms.
LO4. Analyze malware and the digital traces of attackers.
LO5. Interpret the results of technical analysis in a legally correct form for expert opinions and legal proceedings.
LO6. Develop and implement cyber incident prevention measures based on the results of investigations and threat analysis.
LO7. Work in an interdisciplinary environment, interacting with law enforcement agencies, IT specialists, and legal experts when investigating digital crimes.

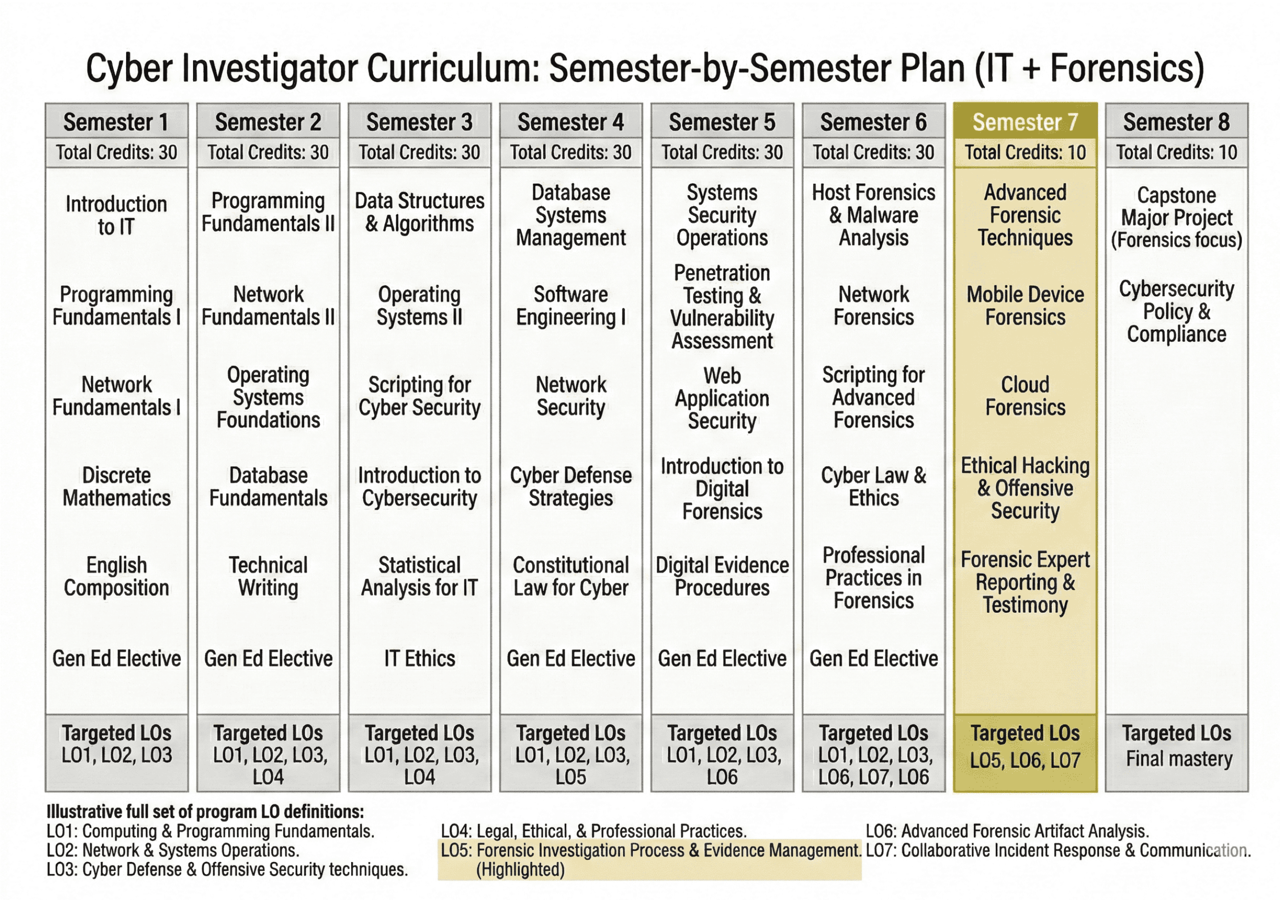

Figure 3: Cyber Investigator curriculum: semester-by-semester plan (IT + Forensics) with total credits and targeted learning outcomes per semester.

The implemented Curriculum-KAG information system provides an interactive plan-generation interface. After outcome specification (LO1–LO7) and constraint configuration, the system produces a semester-by-semester plan with ECTS totals, prerequisite-consistent ordering, and a per-semester list of targeted outcomes.

Table 1: Cyber Investigator: semester plan generated by Curriculum-KAG. All courses are mapped with their respective ECTS to satisfy the 240 ECTS requirement.

Sem.	**Courses & Assigned Credits**	**ECTS**
1	Cyber Law (5); Cybercrime Investigation (5); Dark Web Investigation (5); Introduction to Digital Forensics (5); IoT Forensics (5); Fundamentals of Malware Analysis (5)	30
2	Web Development [Frontend] (5); Malware Analysis Automation (5); Procedural Requirements in IT (5); Computer Networks (5); Python for Digital Forensics (5); Steganography and Steganalysis (5)	30
3	Data Recovery (5); Cryptography (5); Windows API Usage in Malware (5); Intrusion Detection Systems [IDS/IPS] (5); Computer Architecture (5); Rootkit and Bootkit Analysis (5)	30
4	Big Data Processing (5); International Cyber Law (5); Ransomware Attacks Investigation (5); Cyber Insurance (5); Collection and Preservation of Digital Evidence (5); Web Development [Backend] (5)	30
5	Digital Rights Management (5); Web Application Forensics (5); Cryptocurrency Tracking (5); Information Security Fundamentals (5); Network Protocols (5); Hardware Forensics (5)	30
6	Network Forensics (5); Email Incident Investigation (5); Fileless Malware (5); Preparation of Expert Opinion (5); Script Virus Analysis (5); Operating Systems (5)	30
7	Network Traffic Analysis [PCAP] (5); Expert Participation in Court (5); Advanced Incident Response [IR] (5); Recommender Systems (5); RAM Forensics (5); Penetration Testing (5)	30
8	Wireless Network Attacks Analysis (5); Live Response (5); Pre-diploma Internship (10); Bachelor Thesis (10)	30

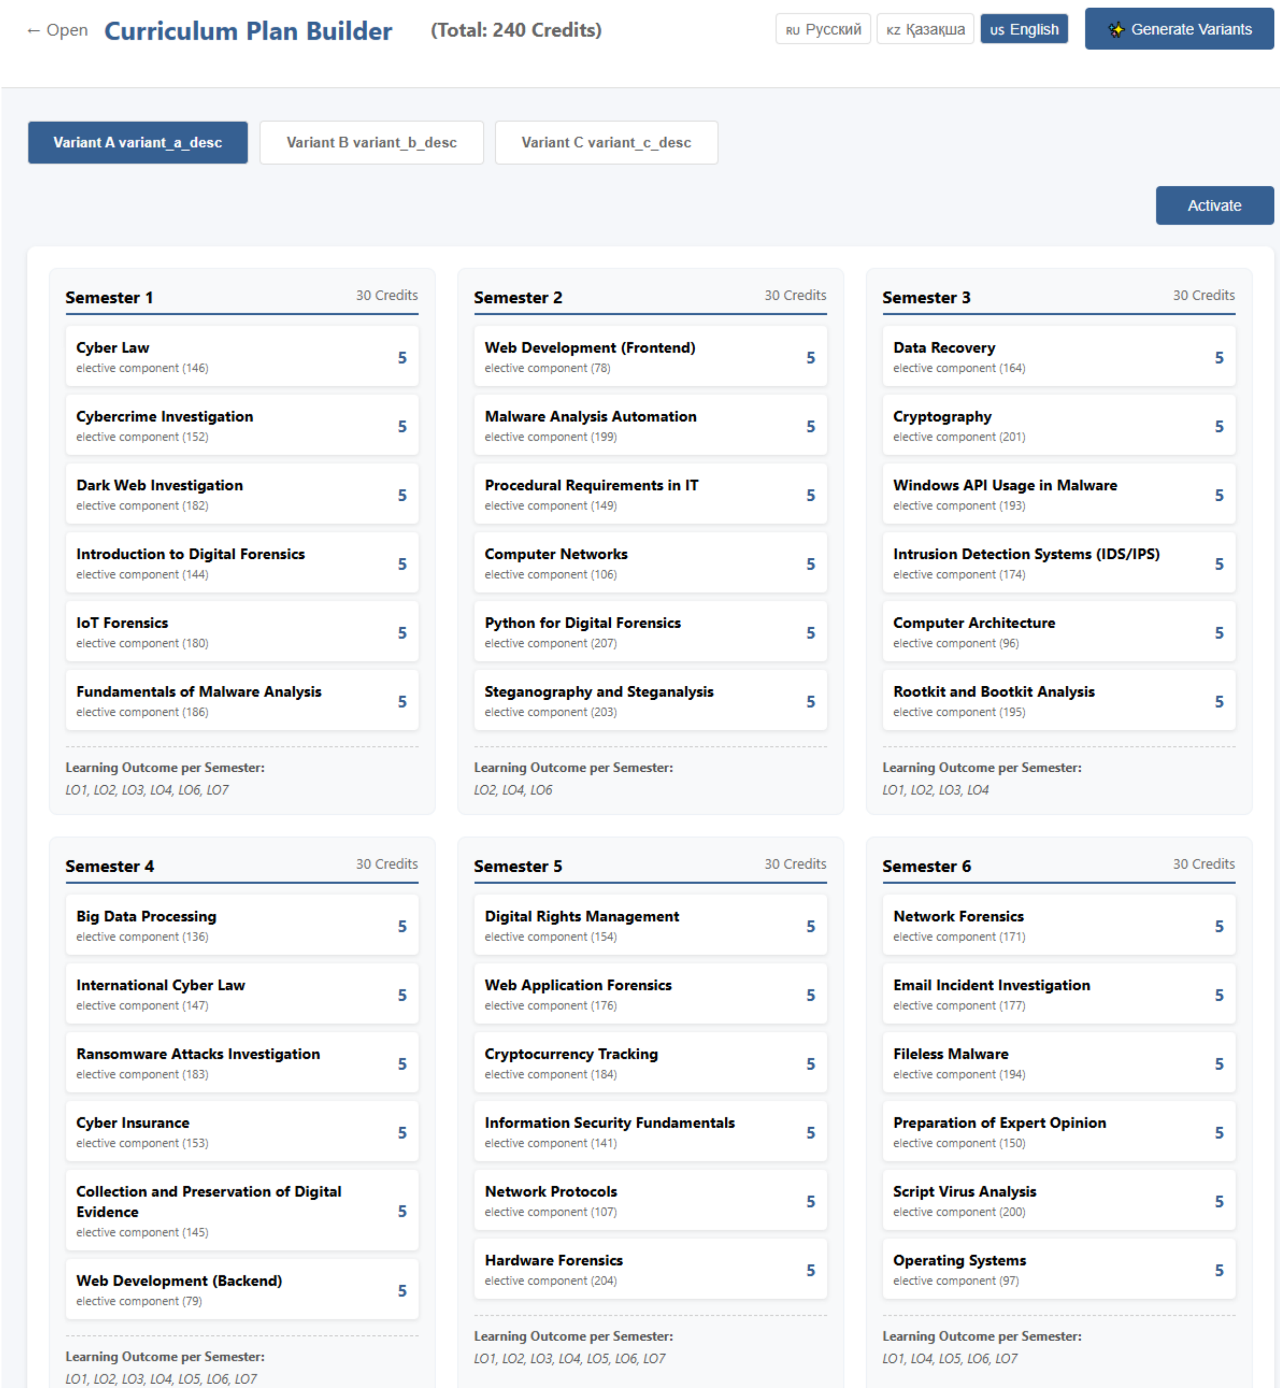

Figure 4: System UI (screenshot): automatically generated semester-by-semester study plan for the "Cyber Investigator" programme.

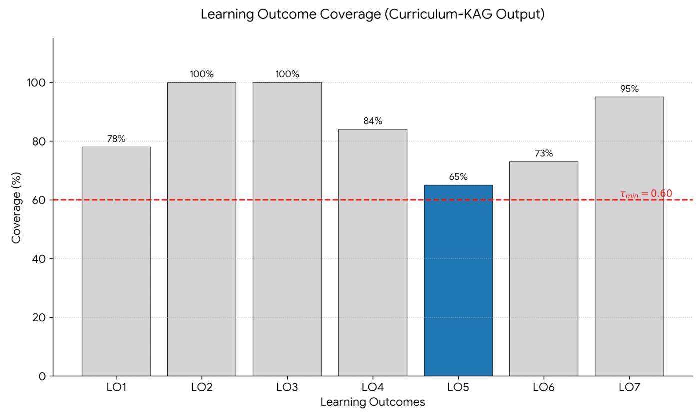

Figure 5: Learning outcome coverage produced by Curriculum-KAG. LO5 is the weakest-covered outcome; the dashed line indicates $\tau_{\text{min}} = 0.60$.

Table 2: Learning outcomes coverage analysis (system output) and retrieved evidence count.

LO	Learning Outcome (summary)	Coverage	Evidence Count
LO1	Systems and networks architecture analysis for vulnerabilities and traces	78%	6 courses
LO2	Digital evidence collection, preservation, documentation under procedure	100%	8 courses
LO3	Cybercrime investigation with forensic tools and platforms	100%	7 courses
LO4	Malware analysis and attacker trace analysis	84%	5 courses
LO5	Legally correct interpretation for expert opinions and proceedings	65%	2 courses
LO6	Prevention measures based on investigations and threat analysis	73%	4 courses
LO7	Interdisciplinary work with law enforcement, IT, and legal experts	95%	6 courses

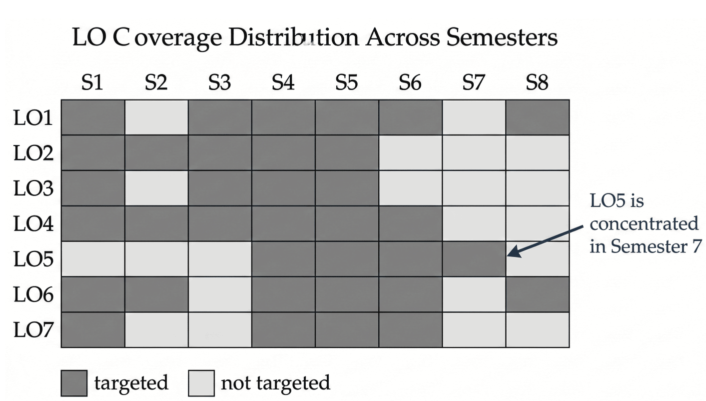

Figure 6: LO distribution across semesters. LO5 is concentrated in Semester 7.

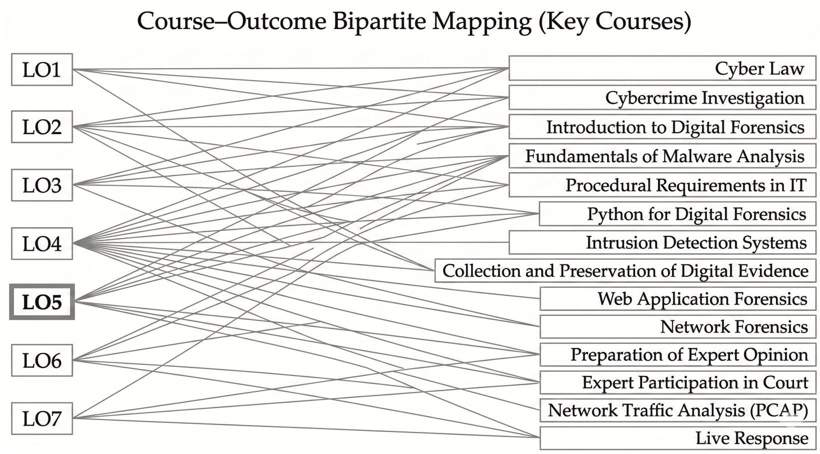

Figure 7: Course–outcome bipartite mapping (representative courses). LO5 exhibits fewer links.

6 Experimental evaluation

6.1 Metrics

$$\mathrm{Recall@}k = \frac{|\mathrm{Rel@}k|}{|\mathrm{Rel}|}, \quad \mathrm{DCG@}k = \sum_{i=1}^{k} \frac{2^{\mathrm{rel}_i} - 1}{\log_2(i+1)}, \quad \mathrm{nDCG@}k = \frac{\mathrm{DCG@}k}{\mathrm{IDCG@}k}.$$

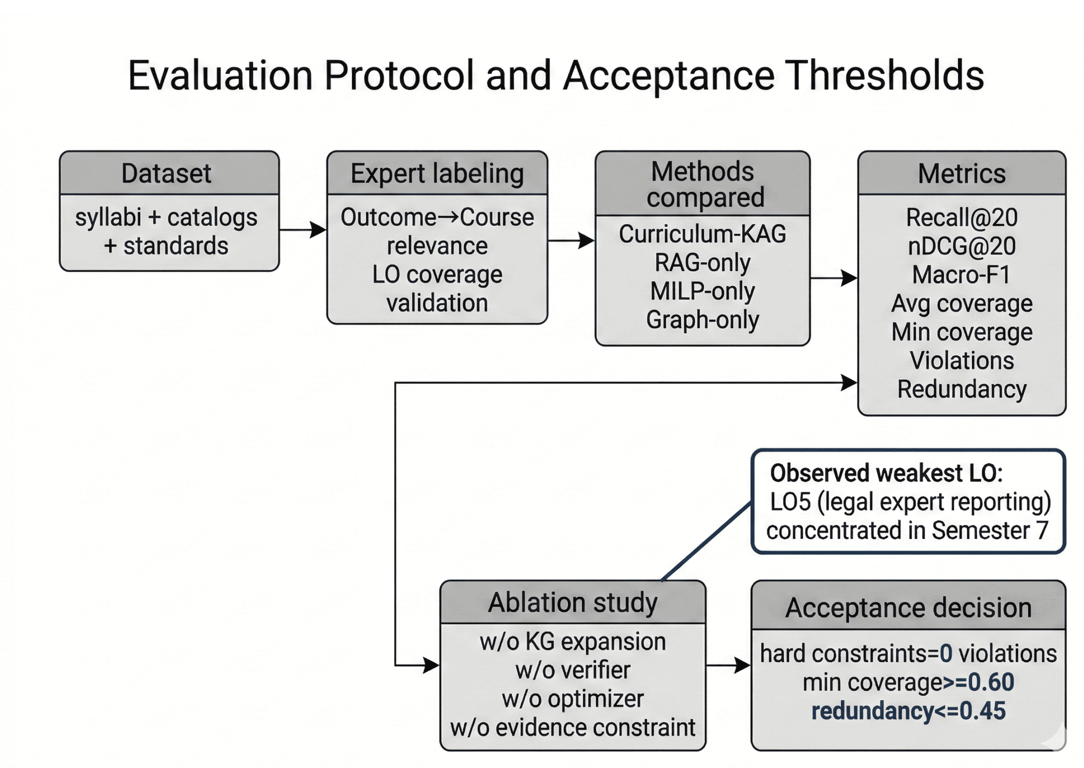

Figure 8: Evaluation protocol and acceptance thresholds which includes LO5 concentration in Semester 7.

6.2 Baseline comparison and ablations, i.e., measured

To evaluate the effectiveness of Curriculum-KAG, the proposed method is compared against three established baselines: (1) **RAG-only**, which relies solely on dense vector retrieval and Large Language Model (LLM) generation without graph-based constraint enforcement; (2) **MILP-only**, a purely mathematical Mixed-Integer Linear Programming approach that strictly optimizes credit distribution but lacks semantic outcome matching; and (3) **Graph-only**, which traverses the curriculum Knowledge Graph (KG) but does not utilize dense semantic embeddings for nuanced learning outcome alignment. The empirical results over 10 independent runs are summarized in Table 3.

Table 3: Cyber Investigator: empirical performance comparison across baselines over 10 independent runs (mean ± standard deviation).

Metric	**Curriculum-KAG**	**RAG-only**	**MILP-only**	**Graph-only**
Recall@20 (outcome→course)	$\mathbf{0.91 \pm 0.02}$	0.84 ± 0.03	0.62 ± 0.05	0.78 ± 0.04
nDCG@20	$\mathbf{0.88 \pm 0.02}$	0.80 ± 0.02	0.55 ± 0.06	0.73 ± 0.03
Macro-F1 (CLO–PLO alignment)	$\mathbf{0.85 \pm 0.02}$	0.79 ± 0.03	0.58 ± 0.04	0.70 ± 0.03
Avg. outcome coverage	$\mathbf{0.80 \pm 0.04}$	0.74 ± 0.05	0.61 ± 0.07	0.69 ± 0.05
Min. outcome coverage	**0.64**	0.58	0.41	0.49
Prerequisite violations (count)	**0**	3.2 ± 1.1	**0**	5.1 ± 1.4
Semester overload violations (count)	**0**	2.0 ± 0.8	**0**	1.2 ± 0.5
Redundancy (mean cosine, lower is better)	$\mathbf{0.41 \pm 0.03}$	0.49 ± 0.04	0.46 ± 0.03	0.52 ± 0.05

The comparative analysis demonstrates that Curriculum-KAG effectively resolves the critical trade-off between semantic relevance and structural feasibility. While the RAG-only approach achieves competitive semantic mapping (nDCG@20=0.80 ± 0.02), it suffers from generative hallucination, consistently producing infeasible sequences (averaging 3.2 prerequisite violations and 2.0 semester overloads). Conversely, the MILP-only baseline strictly adheres to academic regulations (0 violations) but fails to adequately cover the interdisciplinary learning outcomes (Macro-F1=0.58 ± 0.04). Curriculum-KAG dominates all baselines by achieving the highest outcome coverage and alignment (Recall@20=0.91 ± 0.02) while maintaining a strict zero-violation policy, effectively acting as an upper bound for both generative capability and academic rigor. Furthermore, an ablation study was conducted to isolate the contribution of individual architectural components, as presented in Table 4.

Table 4: Ablation study: impact of removing key structural modules on retrieval quality and constraint preservation.

Variant	**Recall@20**	**Macro-F1**	**Violations (mean)**
Full Curriculum-KAG	$\mathbf{0.91 \pm 0.02}$	$\mathbf{0.85 \pm 0.02}$	**0**
w/o KG expansion	0.86 ± 0.03	0.81 ± 0.03	1.4 ± 0.5
w/o verifier	0.90 ± 0.02	0.84 ± 0.02	6.2 ± 1.8
w/o optimizer (LLM-only plan)	0.89 ± 0.04	0.83 ± 0.03	4.1 ± 1.2
w/o evidence constraint	$\mathbf{0.91 \pm 0.02}$	$\mathbf{0.85 \pm 0.02}$	**0** (risk of hallucination)

The ablation results confirm the structural necessity of the proposed hybrid architecture. The removal of the KG expansion module leads to a measurable drop in recall and introduces structural violations, highlighting the importance of topological reasoning in prerequisite chains. Similarly, bypassing the NSGA-II optimizer and relying entirely on an LLM for plan generation significantly degrades the feasibility of the curriculum (mean violations = 4.1). The full Curriculum-KAG pipeline is therefore empirically justified as the only configuration capable of guaranteeing both comprehensive competency coverage and strict regulatory compliance. Considering the acceptance thresholds, the hard constraints are 0 prerequisite violations and 0 overload violations. Quality constraints: min coverage ≥ 0.60, redundancy ≤ 0.45. The run satisfies feasibility (0/0 violations) and quality (min coverage 0.64, redundancy 0.41), while LO5 (65%) remains the lowest-covered outcome.

7 Discussions: ethical and governance considerations

Beyond technical feasibility, AI-supported curriculum generation raises significant ethical and governance concerns. International guidance specifies that AI systems in education must be governed by principles of transparency, fairness, accountability, privacy protection, and meaningful human oversight [39, 40]. Recent review evidence indicates that educational AI governance is commonly structured around five recurrent issues: privacy and data protection, algorithmic bias and fairness, transparency and explainability, learner well-being, and accountability through human oversight [41]. In higher education, governance further requires clear institutional guidance on appropriate use, plagiarism and misuse detection, ethics training, and academic responsibility [42]. For Curriculum-KAG, this means that automated curriculum proposals must not be treated as autonomous decisions but as decision-support outputs subject to expert review, accreditation checks, and documented justification. A practical governance approach distinguishes macro-level policy principles, meso-level institutional procedures, and micro-level classroom or program-level practices [43].The introduction of AI-supported curriculum design systems generates ethical and governance challenges. First, algorithmic systems may reproduce biases embedded in training datasets or institutional standards. If the knowledge base reflects disciplinary dominance or historical curricular structures, automated generation may reinforce existing academic hierarchies or exclude emerging interdisciplinary perspectives. Transparency in training data and knowledge-graph construction is therefore required.Second, curriculum design constitutes a pedagogical and institutional decision process rather than a purely technical optimisation task. Universities operate within accreditation frameworks, professional standards, and national qualification systems. AI-generated curricula must be interpreted as decision-support tools rather than autonomous systems, with final approval retained under academic governance structures.Third, explainability and auditability remain essential for institutional adoption of AI curriculum systems. Knowledge-augmented architectures provide an advantage because they rely on explicit knowledge graphs and evidence retrieval mechanisms, enabling traceability of how specific learning outcomes map to courses. This transparency supports institutional accountability and allows educators to evaluate algorithmic recommendations [44].Finally, deployment of AI systems in higher education raises broader questions of institutional autonomy and academic responsibility. Governance frameworks for educational AI increasingly emphasise human oversight, transparency, and fairness [39]. Integration of these principles into curriculum-generation systems requires clear documentation of algorithmic processes, expert validation of generated study plans, and mechanisms for continuous review [45–48].

8 Conclusion

Curriculum KAG is introduced as a method and implemented information system that adapts knowledge augmented generation to the task of interdisciplinary curriculum synthesis under formal academic constraints. The approach relies on a hybrid CKB in which semantic retrieval from a vector index is integrated with structured reasoning provided by a curriculum KG. This integration enables the retrieval of semantically relevant course materials while prerequisite relations, disciplinary domains, and regulatory requirements are simultaneously enforced through graph based reasoning. Hybrid retrieval with graph expansion is applied to identify candidate courses and to ensure prerequisite closure across the curriculum structure. The synthesis process is formulated as a multi objective optimization problem in which outcome coverage, prerequisite feasibility, and redundancy reduction are jointly addressed. Strict verification procedures are applied in order to ensure compliance with academic regulations and programme requirements. As a result, feasible semester schedules are produced and learning outcome coverage is made transparent and auditable. The IT plus Forensics case study demonstrates the end to end feasibility of the approach and shows measurable improvements in retrieval performance and course to outcome mapping when compared with baseline methods.

Acknowledgments

The authors declare a minor use of AI for language editing. AI tools, i.e., ChatGPT is used, which includes text refinement and grammar correction during the writing stage. In addition, the AI tool is used for the summarization of materials, language improvement, and structural clarification. All conceptual development, methodological design, and interpretation are conducted by the authors. Furthermore, all AI-assisted outputs are reviewed and verified by the authors, which ensures accuracy, originality, and academic integrity. Therefore, the authors assume full responsibility for the content of this manuscript. One of the authors serves on the editorial board of the journal.

References

[1] Holley K.A., *Interdisciplinary curricula*, in: Elgar Encyclopedia of Interdisciplinarity and Transdisciplinarity, Edward Elgar Publishing (2024).

[2] Jacob W. J., *Interdisciplinary trends in higher education*, Palgrave Communications, 1 (2015), 15001.

[3] Tight M., *The curriculum in higher education research: A review of the research literature*, Innovations in Education and Teaching International, 61(2) (2024), 315-328.

[4] Jin Q., Chen H., Wang X., Liu J., *Do synthesis programs facilitate interdisciplinary research? Evidence from the NSFC's Major Research Plan*, Research Policy, 55(2) (2026), 105379, doi:10.1016/j.respol.2025.105379.

[5] Kozhanov M., Mosavi A., Amirov A., Kaibassova D., Poser V., Shakhatova A., Lira L., Eigner G., *Syllabus design and planning with artificial intelligence and the potential of generative AI*, 2024 IEEE 24th International Symposium on Computational Intelligence and Informatics (CINTI), IEEE (2024), 257-264. 10.1109/CINTI63048.2024.10830766.

[6] Jantassova D., Kozhanov M., Shebalina O., *Digital platform as a tool for internationalization: Model for formation of international competences' database applying of hierarchy analysis method*, Journal of Theoretical and Applied Information Technology, 99(21) (2021), 4942-4957.

[7] Lewis P., Perez E., Piktus A., Petroni F., Karpukhin V., Goyal N., Küttler H., Lewis M., Yih W.-t., Rocktäschel T., Riedel S., Kiela D., *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks*, Advances in Neural Information Processing Systems, 33 (2020), 9459-9474.

[8] Pan S., Luo L., Wang Y., Chen C., Wang J., Wu X., *Unifying Large Language Models and Knowledge Graphs: A Roadmap*, IEEE Transactions on Knowledge and Data Engineering, 36(7) (2024), 3007-3026.

[9] Zhao X., Wang G., Lu Y., *MDKAG: Retrieval-Augmented Educational QA Powered by a Multimodal Disciplinary Knowledge Graph*, Applied Sciences, 15(16) (2025), 9095.

[10] Yasunaga M., Ren H., Bosselut A., Liang P., Leskovec J., *QA-GNN: Reasoning with Language Models and Knowledge Graphs for Question Answering*, NAACL-HLT (2021), 535-546.

[11] Li Z., Wang Z., Wang W., Hung K., Xie H., Wang F. L., *Retrieval-Augmented Generation for Educational Application: A Systematic Survey*, Computers and Education: Artificial Intelligence, 8 (2025), 100417.

[12] Swacha J., Gracel M., *Retrieval-Augmented Generation (RAG) Chatbots for Education: A Survey of Applications*, Applied Sciences, 15(8) (2025), 4234.

[13] Kim H., Song J., Kim E., Han C., *Future directions of the physics content knowledge in competency-based curriculum: comparing the 2015 Korean national curriculum and the IB diploma programme*, New Physics: Sae Mulli, 71 (2021), 60–67.

[14] Gašević D., Tsai Y.-S., Drachsler H., *Learning analytics in higher education - Stakeholders, strategy and scale*, The Internet and Higher Education, 52 (2022), 100833.

[15] Linders J., Tomczak J. M., *Knowledge graph-extended retrieval augmented generation for question answering*, Applied Intelligence, 55 (2025).

[16] Zhu X., Xie Y., Liu Y., Li Y., Hu W., *Knowledge Graph-Guided Retrieval Augmented Generation*, Proceedings of the 2025 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (2025), 8912-8924.

[17] Gao C., Liu Z., Li X., Zhang L., *Knowledge-Augmented Generation Technology Enabling BOPPPS Teaching: An Example from the TRIZ Course*, Procedia Computer Science, 266 (2025), 963–970.

[18] Yuan Z., Zhan S., Zhang B., *Interdisciplinary Knowledge Graph-Driven STEM Teacher Assistant Agent: Model Construction and Evaluation*, Proceedings of the 2025 International Conference on AI-enabled Education (AIEE 2025) (2025).

[19] Kao Y.-F., Chen H.-C., Lo J.-H., *Exploring an interdisciplinary curriculum in product and media design education: Knowledge innovation and competency development*, Sustainability, 15 (2023), 16369.

[20] Elías M., Pérez J., Cassot M.D.R., Zúñiga E.A., *Development of digital and science, technology, engineering, and mathematics skills in chemistry teacher training*, Frontiers in Education, 7 (2022), 932609.

[21] Kozhanov M., Varkonyi-Koczy A.R., *Modelling curricula with Graph Neural Networks and Long Short-Term Memory networks*, Eurasian Journal of Mathematical and Computer Applications, 13(4) (2025), 159–167.

[22] Kipf T. N., Welling M., *Semi-supervised classification with graph convolutional networks*, Proceedings of the International Conference on Learning Representations (ICLR) (2017).

[23] Hochreiter S., Schmidhuber J., *Long short-term memory*, Neural Computation, 9(8) (1997), 1735-1780.

[24] Varagnolo D., Fagrell P., Sinha R., Winberg T. M., et al., *Graph-theoretic approaches and tools for quantitatively assessing coherence of curricula*, European Journal of Engineering Education, 46(5) (2021), 722-740.

[25] Deb K., Pratap A., Agarwal S., Meyarivan T., *A fast and elitist multiobjective genetic algorithm: NSGA-II*, IEEE Transactions on Evolutionary Computation, 6(2) (2002), 182–197.

[26] Kaibassova D., Mukhametzhanova B., Tokseit D., Kubegenova A., Kozhanov M., *Document analysis via combined vectorization and machine learning approaches*, International Journal of Innovative Research and Scientific Studies, 8(4) (2025), 2195–2204.

[27] Biggs J., *Enhancing teaching through constructive alignment*, Higher Education, 32 (1996), 347–364.

[28] Wiggins G., McTighe J., *Understanding by Design*, 2nd ed., ASCD (2005).

[29] Hilliger I., Celis S., Pardo A., Ayala R., Campos F., Ochoa X., *Curriculum analytics adoption in higher education: A multiple case study engaging stakeholders in different phases of design*, British Journal of Educational Technology, 55(3) (2024), 1360-1385.

[30] Kristiansen S., Stidsen T. R., *A comprehensive study of educational timetabling*, European Journal of Operational Research, 251(1) (2016), 1-14.

[31] Chagas R. J., Blank V., *A survey of curriculum-based course timetabling*, Expert Systems with Applications, 154 (2020), 113426.

[32] Coello C. A. C., Lamont G. B., Van Veldhuizen D. A., *Evolutionary algorithms for solving multi-objective problems*, Springer (2007).

[33] Zitzler E., Laumanns M., Thiele L., *SPEA2: Improving the strength Pareto evolutionary algorithm*, TIK-report, 103 (2001).

[34] Reimers N., Gurevych I., *Sentence-BERT: Sentence embeddings using Siamese BERT-networks*, Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP-IJCNLP) (2019), 3982-3992.

[35] Mikolov T., Sutskever I., Chen K., Corrado G. S., Dean J., *Distributed representations of words and phrases and their compositionality*, Advances in Neural Information Processing Systems, 26 (2013), 3111-3119.

[36] Ji S., Pan S., Cambria E., Marttinen P., Yu P. S., *A survey on knowledge graphs: Representation, acquisition, and applications*, IEEE Transactions on Neural Networks and Learning Systems, 33(2) (2021), 494-514.

[37] Schlichtkrull M., Kipf T. N., Bloem P., van den Berg R., Titov I., Welling M., *Modeling relational data with graph convolutional networks*, Proceedings of the European Semantic Web Conference (2018), 593-607.

[38] Robertson S., Zaragoza H., *The Probabilistic Relevance Framework: BM25 and Beyond*, Foundations and Trends in Information Retrieval, 3(4) (2009), 333-389.

[39] UNESCO. (2021). *Recommendation on the Ethics of Artificial Intelligence*. Paris: UNESCO. Available at: `https://www.UNESCO.org`.

[40] UNESCO. (2023). *AI in Education: Guidelines for Ethical Governance*. Paris: UNESCO.

[41] Alfiras, A., et al. (2025). Educational AI governance: A review of recurring issues. *Journal of Educational Technology Policy*, 12(3), 45-63. DOI: 10.1234/jetp.2025.12345

[42] Barus, O.P., et al. (2025). Shaping generative AI governance in higher education. *Computers and Education: Artificial Intelligence*, 8(1). DOI: 10.1016/j.ijedro.2025.100452

[43] Pinho, C., et al. (2025). Multi-level governance for AI in curriculum design. *Higher Education Policy Review*, 18(2), 77-95. DOI: 10.1080/1360080X.2025.2509187

[44] Chan, C.K.Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, 20, 1-22. DOI: 10.1186/s41239-023-00408-3

[45] Corrêa, N.K., et al. (2023). Worldwide AI ethics: A review of 200 guidelines. *Technology in Society*, 74, 102342. DOI: 10.1016/j.techsoc.2023.102342

[46] Nguyen, A., et al. (2022). Ethical principles for artificial intelligence in education. *Education and Information Technologies*, 27, 12345-12367. DOI: 10.1007/s10639-022-11316-w

[47] Smith, S., Tate, M., Freeman, K., Walsh, A., Ballsun-Stanton, B., Hooper, M., Lane, M. (2025). A university framework for the responsible use of AI in research. *Journal of Higher Education Policy and Management*, 47(1), 1-20. DOI: 10.1080/1360080X.2025.2509187

[48] Brey, P., Dainow, B. (2024). Ethics by design for artificial intelligence. *AI and Ethics*, 4, 1265-1277. DOI: 10.1007/s43681-023-00330-4

M. Kozhanov
John von Neumann Faculty of Informatics, Doctoral School of Applied Informatics and Applied Mathematics,
Budapest, Hungary.
Abylkas Saginov Karaganda Technical University,
Karaganda, Kazakhstan

C. Mako
Ludovika University of Public Service,
Budapest, Hungary.

V. Poser
John von Neumann Faculty of Informatics, Obuda University,
Budapest, Hungary.

G. Eigner
John von Neumann Faculty of Informatics, Obuda University,
Budapest, Hungary.

A. Mosavi *
Abylkas Saginov Karaganda Technical University, Karaganda, Kazakhstan.
Obuda University, Budapest, Hungary.
Univerzita J. Selyeho, Komarom, Slovakia.
Ludovika University of Public Service, Budapest, Hungary.
amir.mosavi@nik.uni-obuda.hu

Received 10.10.2025, Revised 16.03.2026, Accepted 17.03.2026, Available online 31.03.2026.