scientific reports

OPEN

Adaptive course recommendation using federated learning and graph convolutional networks in IoT-enhanced e-learning

Huizhong Pu[1]✉ & Yan Hua[2]

The increase in e-learning platforms, especially Massive Open Online Courses (MOOCs), highlights the necessity for sophisticated, privacy-conscious recommendation algorithms that adjust to evolving learner interactions in IoT-integrated settings. This study introduces an innovative architecture that utilizes Federated Learning (FL) to safeguard user privacy during distributed training on educational platforms. This approach utilizes Graph Convolutional Networks (GCN) to depict intricate user-course interactions as a graph, adeptly capturing higher-order relational dependencies. Furthermore, DistilBERT-based feature extraction generates concise, semantically dense representations from course descriptions, hence improving content relevancy. Real-time IoT data, including user engagement metrics from smart devices, dynamically influences graph connections, facilitating context-aware recommendations.The suggested solution emphasizes scalability and privacy, tackling essential issues in contemporary e-learning environments. Thorough assessments indicate that our methodology substantially surpasses baseline methodologies across various performance indicators, providing exceptionally tailored course recommendations. This research promotes the advancement of adaptive, safe, and efficient recommendation systems for IoT-integrated e-learning, enhancing engaging and personalized learning experiences for users globally.

Keywords E-learning, Course recommendation, Federated learning (FL), Graph convolutional networks (GCN), DistilBERT, IoT, Privacy-preserving, Personalized learning, Scalability, MOOCs

The swift proliferation of e-learning platforms, especially MOOCs, has transformed global education by providing a variety of courses across disciplines to millions of learners globally[1]. This digital transition fosters lifelong learning but presents considerable challenges in providing personalized course recommendations due to the extensive diversity of courses and heterogeneous learner profiles[2]. The incorporation of Internet of Things (IoT) technology into e-learning systems improves these platforms by collecting real-time engagement metrics, including clicks, duration of interaction, and sensor-driven activities, facilitating dynamic and context-sensitive learning experiences[3]. Recommendation systems are essential to addressing these challenges by suggesting courses tailored to individual preferences[4].

Conventional collaborative filtering techniques rely on user ratings and enrollment behaviors to identify similarities between users or courses[5]. Recent studies have examined content-based filtering, employing course metadata such textual descriptions to suggest analogous learning resources[6]. Although these strategies enhance content relevance, they frequently neglect data privacy and real-time flexibility, both of which are essential in IoT-augmented e-learning systems. Deep learning methodologies have also been employed to integrate contextual elements and learner attributes into recommendation models[7]. Trust-based methodologies have employed clustering techniques to improve the reliability of recommendations[8]. Nonetheless, centralized data processing in numerous systems presents considerable privacy issues, especially in scattered IoT-integrated settings[9].

Personalized recommendation systems are crucial for enhancing learner engagement, course completion rates, and educational results by tailoring suggestions to specific learning requirements[10]. IoT data provide real-time adjustments, guaranteeing that recommendations stay pertinent in evolving situations[11]. Privacy-preserving methodologies, such FL, comply with international data protection laws such as GDPR, rendering them essential for secure e-learning frameworks[12]. Notwithstanding these gains, current methodologies encounter significant constraints. Collaborative filtering techniques fail to adequately represent intricate, higher-order user-course

[1]School of Accountancy, Wuxi City College of Vocational Technology, Wuxi 214153, Jiangsu, China. [2]School of Design, Wuxi Normal College, Wuxi 214153, Jiangsu, China. ✉email: snoopy_phz@outlook.com

interactions, resulting in diminished recommendation accuracy[5]. Centralized training methodologies need the consolidation of sensitive user data, hence presenting privacy problems in IoT contexts[13]. Content-based techniques, however proficient with metadata, frequently neglect to include dynamic engagement patterns, hence constraining flexibility[14]. Deep learning models require substantial computational resources, impeding scalability on large-scale platforms[15]. Trust-based methodologies depend on rudimentary clustering, exhibiting a deficiency in resilience across varied educational contexts[8]. The incorporation of hybrid models that merge collaborative and content-based strategies is still insufficient in effectively tackling both scalability and personalization[16]. The absence of effective privacy-preserving methods in numerous systems intensifies these issues, especially in dispersed situations[17]. Recent investigations have emphasized the necessity for sophisticated graph-based methodologies to represent intricate relationships; however, their utilization in e-learning remains insufficiently investigated[18]. The promise of natural language processing to improve semantic comprehension in recommendations has not been fully actualized[19]. These deficiencies highlight the necessity for an innovative framework that amalgamates sophisticated modeling with privacy-preserving and scalable solutions[20].

This research is motivated by the need to overcome existing restrictions and capitalize on new potential in e-learning. The integration of IoT has the capacity to improve the relevance of recommendations via real-time interaction data; yet, it heightens privacy concerns due to the decentralized creation of data[11].

These difficulties parallel those in federated healthcare systems, where adversarial threats and data-poisoning hazards are paramount for managing sensitive personal information[21,22]. Current methods, including collaborative filtering and deep learning, struggle to balance personalization, privacy, and scalability[5]. Graph-based modeling shows promise in capturing complex relational structures; however, its use in e-learning recommendation systems remains limited[14]. FL has recently emerged as an effective paradigm for distributed decision-making in IoT environments, providing privacy-preserving and communication-efficient coordination among decentralized clients[23]. FL offers a method for privacy-preserving training; nevertheless, its incorporation with graph-based models remains insufficiently investigated[12]. Recent breakthroughs in natural language processing have revealed the capacity to extract significant elements from textual material; yet, their application in e-learning recommendations remains in its infancy[15]. The integration of these technologies presents a distinctive potential to create a resilient, adaptive, and secure recommendation system for contemporary e-learning environments.

The primary aim of this research is to develop an adaptive and privacy-preserving course recommendation framework for MOOC-based e-learning environments. The platform incorporates FL for decentralized model training without data aggregation, GCN to elucidate intricate and higher-order user-course interactions, and DistilBERT to derive semantically rich representations from course descriptions. The integration of these technologies is essential, as each element mitigates a specific limitation: traditional GCN-based recommenders necessitate centralized data, jeopardizing user privacy; FL facilitates distributed learning but is deficient in semantic depth; and DistilBERT improves contextual comprehension. The system integrates these components to deliver adaptable, privacy-preserving, and semantically informed recommendations that concurrently tackle customization, scalability, and privacy.

The suggested technique improves personalization for extensive, diverse learner populations characteristic of MOOCs, where preserving user privacy and scalability poses significant challenges. The system generates adaptive, context-aware, and semantically relevant recommendations by integrating these technologies, in accordance with the changing interactions of learners.

This research's principal contributions are summarized as follows:

- A privacy-preserving federated learning framework that facilitates decentralized model training while adhering to contemporary data protection rules.
- A GCN module designed to capture higher-order, context-aware interactions between users and courses, with edge weights dynamically adjusted based on real-time IoT engagement metrics.
- A DistilBERT-based semantic representation system to produce significant embeddings from course descriptions, improving contextual relevance and precision.
- Thorough studies on actual IoT-augmented e-learning datasets exhibiting consistent performance improvements compared to leading federated and graph-based recommendation systems.
- A scalable and reproducible framework that systematically addresses the privacy–scalability–personalization trade-off and provides a foundation for future adaptive e-learning research.

The subsequent sections of this work are structured as follows: Chap. 2 examines pertinent literature on e-learning recommendation systems, highlighting collaborative filtering, deep learning, and privacy-preserving methodologies. Chapter 3 delineates the proposed framework, encompassing the components of FL, GCN, and DistilBERT, along with implementation specifics and evaluation criteria. Chapter 4 delineates experimental designs, outcomes, and baseline comparisons. Chapter 5 examines the findings, ramifications, and limits. Chapter 6 closes the document and delineates prospective study avenues.

Related works

This section examines current developments in e-learning recommendation systems, focusing on graph-based and federated learning methodologies. It encapsulates their principal contributions, constraints, and significance to the proposed framework.

Recent advancements in GCNs have significantly enhanced the capacity of recommender systems to represent intricate relational relationships between users and objects. LightGCN–PKA[24] amalgamates a Light GCN model with a tailored knowledge-aware attention mechanism, proficiently synthesizing user–item and knowledge graphs. This methodology encapsulates nuanced semantics and attains robust outcomes across multiple benchmark datasets. Nonetheless, its constrained scalability and lack of real-time IoT connectivity diminish its relevance in dynamic e-learning contexts. Likewise, ConceptGCN[25] utilizes knowledge graphs and SBERT embeddings to enhance tailored MOOC recommendations, hence augmenting originality and transparency in CourseMapper. However, its reliance on pre-trained transformers and absence of IoT-driven adaptability limit its use in extensive, time-critical environments.f

Federated Learning (FL) has emerged as a pivotal privacy-preserving framework, facilitating decentralized model training without the exchange of raw data. The FedDPRec architecture[26] utilizes dynamic differential privacy to reconcile security and accuracy in news recommendations; nonetheless, it incurs computational complexity and is deficient in semantic modeling. In[27], transformer-based federated learning systems (BERT and BST) enhance accuracy and user behavior modeling, however they are computationally intensive and IoT-agnostic. Incentive-based federated learning frameworks for mobile edge computing (MEC) promote user engagement via reward systems, enhancing scalability while lacking integration of semantic and real-time context.

Federated graph-based solutions have been implemented in extensive IoT applications. For instance, Green-IoT FL[29] illustrates synchronized real-time decision-making across decentralized sensors in safety-critical settings. Despite their efficacy in anomaly identification, such models have not been tailored for educational settings. In FGC[30], a federated GCN-based trust recommendation system maintains anonymity; nevertheless, it employs static graphs, which constrains adaptability. Other experiments, such GFed-PP[31] and FedGA[32], improve aggregation efficiency and personalization in non-IID contexts but remain deficient in semantic depth and IoT-driven responsiveness. Alternative privacy-preserving methodologies, such as distributed data cooperation[33] and dual-cloud models[34], provide privacy protection yet lack the integration of federated graph designs or real-time semantic enrichment, hence diminishing their relevance to dynamic MOOCs. Table 1 delineates exemplary FL–GCN and privacy-conscious recommendation frameworks, detailing their models, datasets, evaluation measures, and principal limitations.

In contrast to current FL–GCN recommendation frameworks that mainly emphasize data privacy preservation or static modeling of user–item connections, the proposed approach presents numerous significant advancements. Initially, it incorporates real-time IoT interaction data to dynamically modify graph edges, facilitating adaptive, context-sensitive recommendations that progress alongside learner behavior. Secondly, it integrates DistilBERT-derived semantic representations of course content, offering a more profound semantic comprehension that conventional GCN-based or federated recommenders do not possess. Third, the framework accomplishes a tri-dimensional optimization of personalization, privacy, and scalability—an element not collectively addressed in previous FL–GCN systems. These modifications provide a privacy-preserving and semantically enriched recommendation process, representing a notable improvement over leading federated graph-based algorithms and providing a personalized learning framework designed for extensive MOOC contexts.

Ref	Year	Approach/model	Dataset(s)	Evaluation metrics	Key limitations
24	2025	LightGCN with personalized knowledge-aware attention	Book-crossing (Book), MovieLens-20 M (Movie), Last.FM (Music) and Dianping-Food (Restaurant	Recall, F1 – score	Lightweight and personalized GCN model; lacks federated setting and real-time IoT adaptation.
25	2024	ConceptGCN with SBERT-based semantic embedding	MOOC & conceptnet datasets	Precision@K, NDCG, accuracy	Strong semantic representation but no federated or privacy-preserving mechanism.
26	2025	FedDPRec: federated learning with differential privacy	E-learning datasets	AUC, MRR, DCG,	Focuses on privacy in FL but ignores graph structure and semantic enrichment.
27	2024	Transformer-based + FL	Movielens 1 m dataset, amazon review dataset	MSE, accuracy, MAE	Improves accuracy in FL; lacks IoT context and graph-aware modeling.
28	2024	FL incentive mechanism for recommendation systems	Industrial datasets	AUC, RMSE	Addresses incentive issues in FL but omits semantic and adaptive graph modeling.
29	2025	FL + greenIOT	Custom dataset: 1900 color images (250 × 250 pixels), 950 Fire and 950 No Fire samples;	Accuracy, loss rate	
30	2023	FGC: GCN-based + FL	Industrial/service data	HR@K, NDCG	Integrates GCN and FL but relies on static graphs without IoT-driven updates.
31	2025	DeFedGCN: decentralized privacy-preserving federated GCN	Benchmark recommendation datasets	RMSE, AUC	Ensures strong privacy but has high computational cost and lacks semantic embeddings.
32	2025	Graph federated learning for personalized privacy recommendation	MovieLens-100 K, MovieLens-1 M, Lastfm-2 K, HetRec2011	HR, NDCG,	Focuses on privacy personalization; lacks semantic representation and IoT adaptivity.
33	2025	Privacy-preserving RS via distributed data collaboration	Simulated distributed datasets	MAE, RMSE	Secure and distributed but not graph-based or semantically enhanced.
34	2024	Privacy-preserving recommendation based on social relationships	FilmTrust (crawled small dataset: 35,497 preference data, 1,853 social data; avg. 24 items/user, 3 friends/user)	Runtime (ms)	Leverages social relations; lacks federated training and graph-based deep modeling.

Table 1. Comparison of recent FL–GCN-based and privacy-aware recommender frameworks.

Attribute	Details
Number of courses	89,000 (explicit interactions)
Number of user interactions	276,000 (implicit interactions)
Metadata	Course titles, textual descriptions, user interaction histories
Interaction types	Ratings, enrollments, clicks, and bookmarks
Source	Mandarine academy platform (2016–2021)

Table 2. Characteristics of the MARS dataset.

Attribute	Details
Number of courses	3,522
Metadata	Duration, language, ratings, instructor names, course titles, descriptions
Interaction types	Enrollments, textual reviews, and learner feedback
Source	Coursera platform (Accessed September 2021)

Table 3. Characteristics of the coursera course review dataset.

Methodology

This chapter delineates the technique of the proposed adaptive course recommendation model, which amalgamates Federated Learning, Graph Convolutional Networks, and DistilBERT inside an Internet of Things-enhanced e-learning environment. The proposed model, in contrast to conventional centralized methods that depend on collaborative filtering and encounter challenges in privacy protection and the modeling of intricate user-course interactions, utilizes decentralized training, graph-based modeling, and semantic embeddings to improve recommendation precision and safeguard user privacy. This chapter delineates the model architecture, the FL framework, implementation specifications, and evaluation criteria, establishing a thorough foundation for the experimental study in the next chapters.

Datasets

This section presents the datasets utilized to assess the proposed Federated GCN–DistilBERT Recommendation Framework. Three actual e-learning datasets were utilized: Mandarine Academy Recommender System (MARS), Coursera Course Review, and Personalized e-Learning Recommendation System. These datasets were chosen for their diversity, the presence of textual metadata, and their capacity to encapsulate semantic, relational, and individualized learner-course interactions. They jointly provide a thorough assessment of the framework's adaptability, scalability, and privacy-preserving features in IoT-integrated systems. Each dataset is elaborated upon in detail below.

Mandarine academy recommender system (MARS)

The MARS dataset offers extensive learner-course interaction data appropriate for assessing personalization and scalability in large-scale online learning contexts. It comprises 89,000 explicit interactions (ratings) and 276,000 implicit interactions (clicks, enrollments, bookmarks), in addition to comprehensive metadata including course titles, descriptions, and user interaction histories. These characteristics facilitate semantic representation learning using DistilBERT and relational modeling via GCNs. Additionally, IoT-driven user engagement measurements are incorporated to dynamically modify edge weights, enhancing the model's contextual responsiveness while preserving privacy via federated learning. Table 2 summarizes the overall properties of the MARS dataset, detailing its scale, information, and interaction aspects.

Coursera course review

The Coursera Course Review dataset, acquired from the Coursera website in September 2021, comprises 3,522 courses along with an extensive array of metadata. This encompasses course duration, language, ratings, instructor details, course titles, descriptions, and enrollment data. The dataset enables content-based and semantic recommendations via DistilBERT embeddings, while timestamped user interactions allow for IoT-driven graph updates within the GCN module. This setting enables the system to deliver adaptive and privacy-preserving recommendations instantaneously. The comprehensive structure and principal attributes of this dataset are delineated in Table 3.

Personalized elearning recommendation system

The Personalized e-Learning dataset emphasizes customized learner models. It consists of 2,000 user profiles and 8,000 learner-course interactions, encompassing metadata that delineates learning styles, academic interests, and participation histories. This information allows the system to customize recommendations based on individual learner preferences and behavioral variability. This is especially appropriate for evaluating the federated learning (FL) aspect, given the data is privacy-sensitive and dispersed across various clients. The system integrates DistilBERT for semantic encoding and GCN for relational learning, effectively capturing both textual and structural aspects of learner behavior. Table 4 summarizes the makeup and properties of the dataset.

Attribute	Details
Number of user profiles	2,000
Number of interactions	8,000
Metadata	Learning styles, academic interests, and engagement histories
Interaction types	Course enrollments, user preferences
Source	Institutional e-learning platform

Table 4. Personalized elearning recommendation system dataset characteristics.

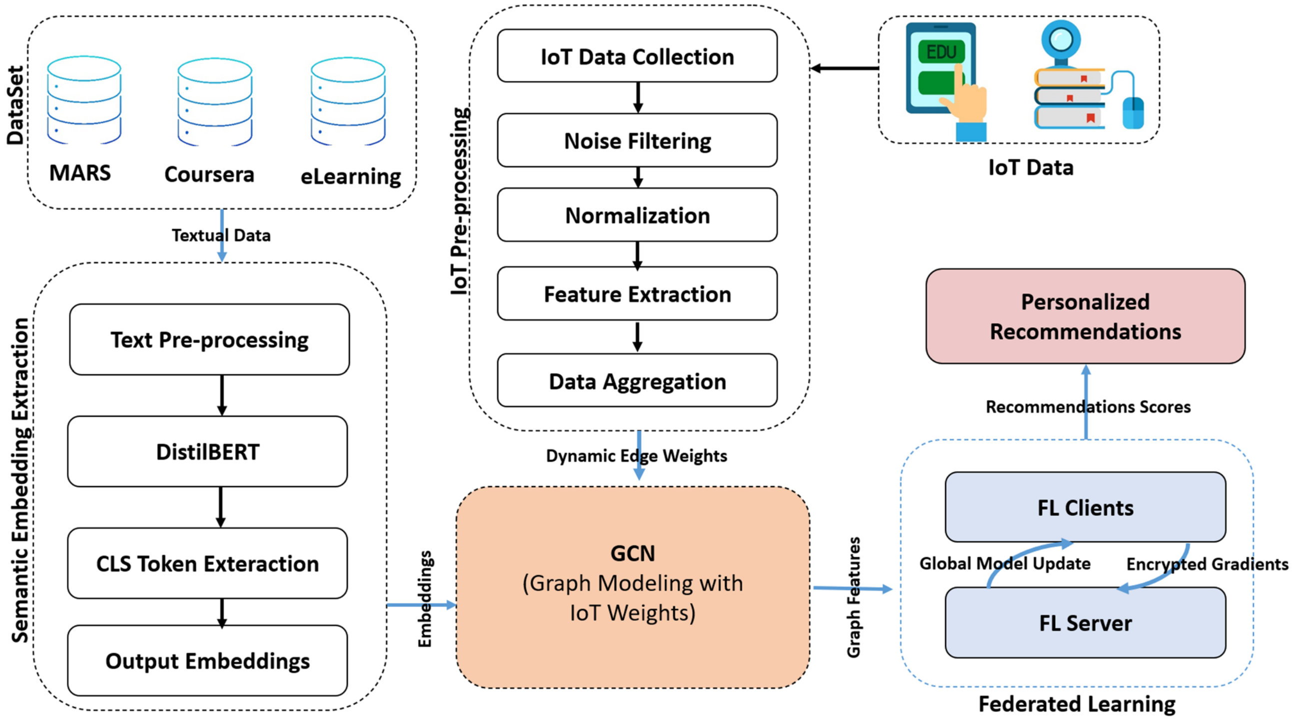

Fig. 1. Proposed recommendations framework.

IoT data
The framework integrates real-time IoT engagement data alongside the three primary datasets to improve the adaptability of the graph-based model. The IoT features encompass click frequency, session duration, device kind, and sensor-based activity metrics, all collected from IoT-enabled educational settings. Each signal facilitates dynamic edge-weight modifications between user and course nodes, enabling the model to represent changing learner engagement patterns. Raw IoT data undergo local processing, filtering, and normalization prior to aggregation within the federated architecture, hence assuring complete privacy preservation.

Proposed model architecture
The design of the proposed Federated GCN–DistilBERT Recommendation Framework, seen in Fig. 1, facilitates the smooth integration of diverse data sources and sophisticated machine learning methodologies to provide adaptive and tailored course suggestions. The model integrates data from three benchmark datasets—MARS, Coursera Course Review, and Personalized e-Learning Recommendation System—with real-time IoT engagement data that records learner interaction parameters, including click frequency, session duration, and device context. This integration enables the system to acquire both semantic and relational dependencies while safeguarding user privacy via decentralized learning.

The overall architecture is composed of four major components:

- Semantic Embedding Extraction using DistilBERT, which transforms course descriptions and user profiles into high-dimensional contextual embeddings.
- Graph-Based Modeling via GCN, which represents user–course relationships and IoT-based interactions as dynamic graphs.
- FL Framework, which enables decentralized model training without sharing raw user data.
- Recommendation Generation, which synthesizes semantic, structural, and contextual information to produce personalized course suggestions.

Each of these components is discussed in detail in the following subsections.

Role and pre-processing of IoT data

The IoT layer is essential for data collecting, local processing, and communication inside the federated architecture. IoT technologies, including wearable sensors, smart classroom elements, and mobile applications, record real-time indications of learner-content interaction, such as session duration, click frequency, and attention metrics. Each IoT node processes this data locally before to its integration into the learner-course interaction graph, so assuring privacy preservation and data dependability. Only pre-processed and aggregated engagement metrics are sent to the client-side GCN module, whereas raw sensor data is retained locally on the device.

The pre-processing of IoT data for GCN edge weighting aims to guarantee data quality, stability, and interpretability prior to its incorporation into the recommendation model. It comprises four mathematical phases: Noise filtering involves the processing of raw IoT signals $S_{\text{raw}}(u, c, t)$, which denote the unprocessed sensor data for user (u), course (c), and time (t), to eliminate outliers.

$$S_{\text{filtered}}(u, c, t) = f_{\text{filter}}(S_{\text{raw}}(u, c, t)) = S_{\text{raw}}(u, c, t) \cdot I\left(\left|S_{\text{raw}}(u, c, t) - \mu_{t-1}\right| < \sigma_{t-1}\right) \tag{1}$$

μ_{t-1} represents the mean of previous readings up to time (t-1), while σ_{t-1} denotes the standard deviation of those readings, serving as statistical measures to eliminate abnormal or noisy sensor outputs. Normalization involves scaling filtered data to the range [0, 1] to maintain consistency across diverse sensors.

$$S_{\text{norm}}(u, c, t) = \frac{S_{\text{filtered}}(u, c, t) - \min(S_{\text{filtered}})}{\max(S_{\text{filtered}}) - \min(S_{\text{filtered}})} \tag{2}$$

where $\min(S_{\text{filtered}})$ and $\max(S_{\text{filtered}})$ denote the minimum and maximum values of the filtered data during the specified time interval; Feature extraction involves calculating engagement indicators as weighted combinations of essential behavioral variables, including average session duration and click frequency.

$$f_{\text{IoT}}(u, c, t) = w_1 \cdot \text{avg\_time}(u, c, t) + w_2 \cdot \text{click\_freq}(u, c, t) \tag{3}$$

where w_1 and w_2 are weighting factors with $w_1 + w_2 = 1$ representing the relative importance of average session time $\text{avg\_time}(u, c, t)$ the normalized average time spent by user (u) on course (c) at time (t) and click frequency $\text{click\_freq}(u, c, t)$ the normalized number of clicks by user (u) on course (c) at time (t); and Data Aggregation, where recent IoT features are aggregated over a temporal window (N) the number of time steps considered using an exponential decay factor (λ) a decay rate between 0 and 1 to capture temporal dynamics with

$$f_{\text{agg}}(u, c, t) = \frac{1}{N} \sum_{i=t-N+1}^{t} \lambda^{t-i} \cdot f_{\text{IoT}}(u, c, i) \tag{4}$$

The aggregated engagement features $f_{\text{agg}}(u, c, t)$ are then integrated into the GCN edge-weight update as

$$w_{u,c}(t) = (1 - \gamma) \cdot w_{u,c}(t-1) + \gamma \cdot f_{\text{agg}}(u, c, t) \cdot I(|\Delta f_{\text{agg}}| > \tau) \tag{5}$$

where $w_{u,c}(t)$ the edge weight between user (u) and course (c) at time (t), (γ) a learning rate between 0 and 1,

$$\Delta f_{\text{agg}} = |f_{\text{agg}}(u, c, t) - f_{\text{agg}}(u, c, t-1)| \tag{6}$$

The absolute change in aggregated features and the sensitivity threshold (τ) dictate the update condition. This multi-stage pre-processing pipeline ensures that only dependable, normalized, and temporally consistent IoT interaction data affect graph updates, thereby enhancing both interpretability and convergence stability of the federated GCN model.

Semantic embedding extraction with distilbert

The DistilBERT module derives semantic embeddings from textual inputs to describe course descriptions and user preferences in a high-dimensional latent space. In contrast to conventional feature engineering methods, DistilBERT automates feature extraction, encapsulating intricate semantic relationships. The procedure illustrated in the "Semantic Embedding Extraction" subgraph of Fig. 1 has four stages:

- Text preprocessing: Course descriptions and user preferences are tokenized with the DistilBERT tokenizer (distilbert-base-uncased). Tokenization transforms text into numerical tokens, applying padding and truncation to maintain consistent sequence lengths (maximum of 512 tokens).
- DistilBERT model: The tokenized inputs are analyzed by the DistilBERT base-uncased model, a streamlined transformer with six layers, 66 million parameters, and a hidden size of 768. The model utilizes self-attention mechanisms to discern contextual linkages within the text.
- CLS token extraction: The CLS token, positioned as the initial token in each sequence, is utilized as a semantic representation of the input text. This 768-dimensional vector embodies the significance of the complete sequence.

- Output embeddings: The CLS token vectors function as 768-dimensional embeddings for courses and users, acting as node features in the ensuing GCN module.

The embeddings offer a strong representation of textual data, allowing the model to identify semantic connections between courses and user preferences.

Graph-based modeling with GCN

The GCN module conceptualizes user-course interactions as a dynamic bipartite graph, with nodes symbolizing users and courses, and edges indicating interactions (e.g., ratings, enrollments, clicks). Figure 2 depicts a bipartite graph, including users U_i on one side and courses C_j on the opposite, interconnected by edges denoting interaction kinds including ratings, enrollments, and clicks. The GCN utilizes tiered convolutional layers to consolidate characteristics from adjacent nodes, harnessing IoT data for real-time edge weighting. Course metadata, analyzed by DistilBERT, augments node features, while user interactions from datasets such as MARS, Coursera Course Review, and Personalized eLearning Recommendation System continuously update the graph, improving recommendation precision and flexibility in an IoT-enhanced e-learning context.

The GCN process, seen in "graph-based modeling" Fig. 3, comprises four steps:

- Graph construction: A bipartite graph $\left(G=(V,\ E)\right)$ is constructed, where $\left(V\right)$ consists of user nodes $\left(U\right)$ and course nodes $\left(C\right)$, and $\left(E\right)$ represents interactions. Course node characteristics are initialized using 768-dimensional DistilBERT embeddings, whilst user nodes utilize embeddings based on user choices or random initialization for new users. Edge weights are initially established based on interaction data (e.g., ratings) and are subsequently modified using IoT metrics (e.g., click frequency, duration of engagement).
- GCN layer 1: The first convolutional layer propagates features using IoT-weighted edges, following the GCN update rule:

$$h_i^{(1)}=\sigma\left(\sum\nolimits_{j\in\mathcal{N}(i)}\frac{e_{ij}}{\sqrt{d_i d_j}}W^{(0)}h_j^{(0)}+b^{(0)}\right) \tag{7}$$

where $\left(h_i^{(1)}\right)$ is the feature vector for node (i), $(\mathcal{N}(i))$is the set of neighbors, (e_{ij}) is the IoT-adjusted edge weight, (d_i) is the node degree, $\left(W^{(0)}\right)$ and $\left(b^{(0)}\right)$ are learnable parameters, and $(\sigma)$is the ReLU activation function. This layer captures local interaction patterns.

- GCN layer 2: The second convolutional layer additionally consolidates features:

$$h_i^{(2)}=\sigma\left(\sum\nolimits_{j\in\mathcal{N}(i)}\frac{e_{ij}}{\sqrt{d_i d_j}}W^{(1)}h_j^{(1)}+b^{(1)}\right) \tag{8}$$

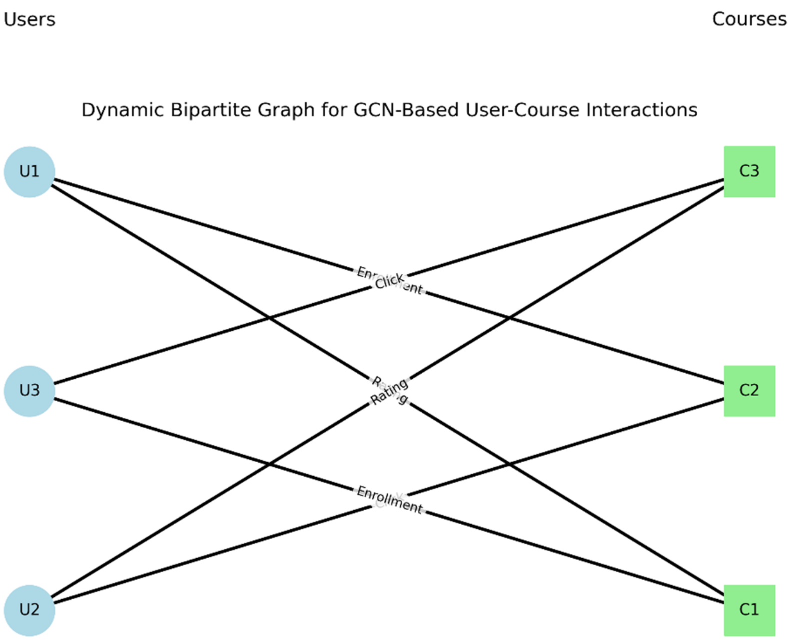

Fig. 2. Dynamic bipartite graph for GCN-base user-course interactions.

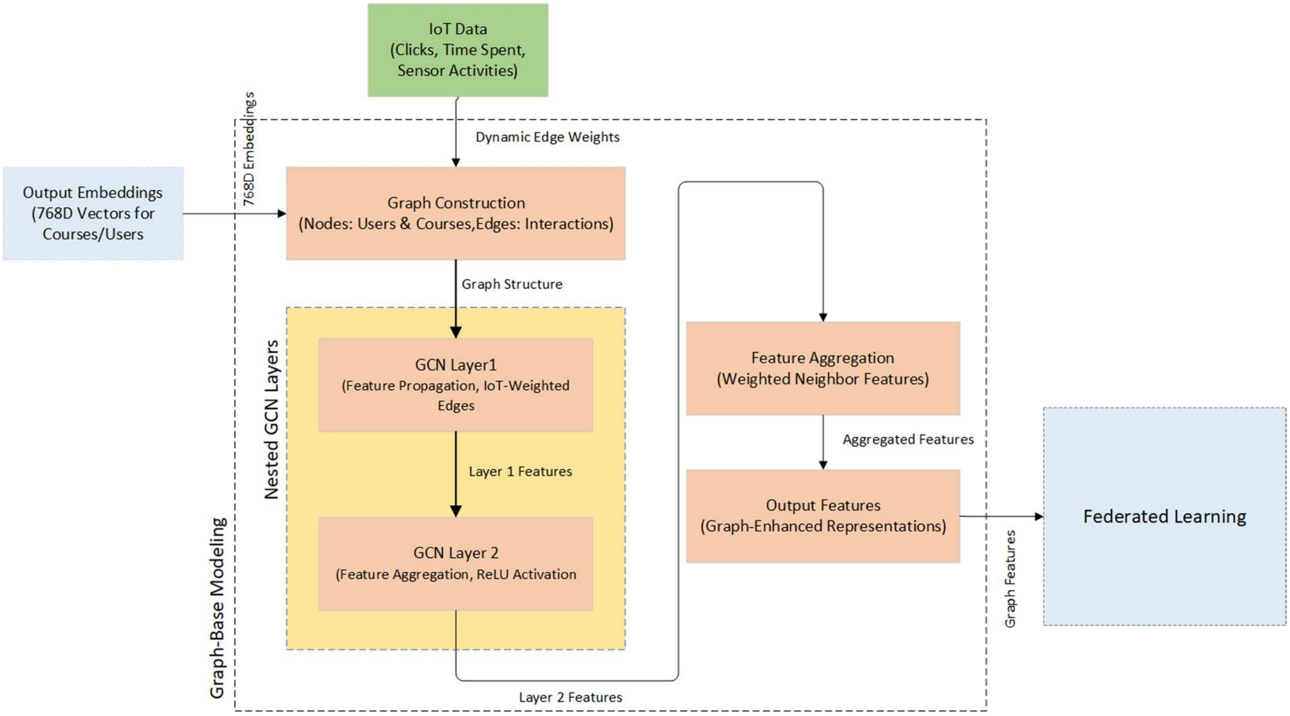

Fig. 3. GCN process in graph-based modeling.

This layer simulates higher-order interactions, augmenting feature representations by ReLU activation to introduce non-linearity.

- Feature aggregation: The consolidated attributes from Layer 2 are integrated to generate enhanced node representations, illustrating both semantic and interaction-based linkages influenced by IoT data.
- Output features: The ultimate GCN output produces graph-enhanced representations for users and courses, utilized for training inside the federated learning framework or for predicting recommendation scores.

The nested GCN layers, shown in the "Nested GCN Layers" sub-subgraph, enable the model to capture multi-hop relational dependencies, significantly improving recommendation accuracy.

IoT-driven edge weight updates in graph connectivity
IoT devices consistently provide engagement signals that dynamically modify the connectivity of the graph-based recommendation model by adjusting the weights of learner-course edges. Each edge $(e_{u,c})$is linked to a time-variable weight $(w_{u,c}(t))$, calculated as a refined amalgamation of historical and real-time engagement data:

$$w_{u,c}(t) = (1-\gamma)\cdot\ w_{u,c}(t-1) + \gamma\ \cdot\ f_{\mathrm{IoT}}(u,c,t)\,, \tag{9}$$

where $(\gamma\ )$ (a smoothing factor ranging from 0 to 1) regulates the pace of temporal adaptation, and $f_{\mathrm{IoT}}(u,c,t)$ denotes the normalized engagement score for u and c at t. Elevated engagement enhances the edge weight, signifying augmented significance, whilst diminished engagement diminishes it, denoting decreased interest. Local clients update these IoT-based weights before each FL communication round, facilitating ongoing adaptation while safeguarding raw IoT data and maintaining privacy. Algorithm 1 enhances this procedure by adaptively modifying edge weights in response to fluctuations in IoT engagement.

To guarantee reproducibility and computational efficiency, edge-weight modifications are executed at predetermined intervals aligned with each federated communication round. In this study, a graph update is planned following each local training epoch, which is about 5 to 10 min of learner engagement. During each period, IoT devices consistently record engagement data; however, updates are consolidated based on events rather than processed in real-time. An edge weight is adjusted just when the cumulative IoT variation $\Delta\ f_{\mathrm{IoT}}$ exceeds a threshold, indicating a substantial shift in engagement. This hybrid approach—integrating interval-based scheduling with event-driven refinement—attains equilibrium among responsiveness to real-time data, model stability, and reproducibility across experimental iterations.

Federated learning framework
The Federated Learning (FL) module in the proposed architecture ensures privacy-preserving training by decentralizing the learning process, thus alleviating privacy concerns linked to centralized methods. The FL architecture, depicted in the "Federated Learning" subgraph in Fig. 1, consists of two primary components: FL

Input:
Previous weights $w_{u,c}(t-1)$
Historical score $r_{u,c}$
IoT signal $f_{IoT}(u,c,t)$
Parameters $(\alpha, \beta, \gamma, \tau)$
Output: *Updated weights* $w_{u,c}(t)$
1. **for each** *edge* $e_{u,c} \in G$ *do*
2. $\Delta f_{IoT} = |f_{IoT}(u,c,t) - f_{IoT}(u,c,t-1)|$
3. *if* $\Delta f_{IoT} > \tau$ *then*
4. $s_{u,c} = \alpha * r_{u,c} + \beta * f_{IoT}(u,c,t)$
5. $w_{u,c}(t) = (1-\gamma) * w_{u,c}(t-1) + \gamma * s_{u,c}$
6. **else**
7. $w_{u,c}(t) = w_{u,c}(t-1)$
8. **end if**
9. **end for**

Algorithm 1. IoT-Driven Edge Weight Update

Clients and FL Central Server, intended to train the GCN-based recommendation engine among decentralized clients while maintaining data privacy. The procedure is designed to optimize computing efficiency, model convergence, and adaptation to diverse data contexts. Each client, generally a user device, trains a local GCN model utilizing its own data, which includes local interactions and IoT measurements. The local loss function, illustrated as mean squared error for rating prediction, is defined as

$$\mathcal{L} = \frac{1}{|E|}\sum\nolimits_{(u,c)\in E}\left(r_{u,c} - \widehat{r}_{u,c}\right)^2 \tag{10}$$

where (ru, c) represents the actual rating and $(\widehat{r}_{u,c})$ denotes the anticipated score for u and c. Local gradients are calculated and transmitted to the server. Clients execute these local model updates following every 10 communication cycles, which is roughly similar to 24 h in a standard MOOC setting. This frequency optimizes computational efficiency and model convergence, enabling enough local training on IoT-augmented graph data prior to synchronization. The update is initiated when local data alterations (e.g., fresh IoT engagement measurements) surpass a threshold, guaranteeing responsiveness to fluctuating user behavior. The FL Central Server consolidates gradients via the FedAvg algorithm, updating the global model parameter $\left(\theta^{(t)}\right)$ as

$$\theta^{(t+1)} = \theta^{(t)} - \eta \sum k = 1^K \frac{n_k}{n} \nabla \mathcal{L}k\left(\theta^{(t)}\right) \tag{11}$$

where $\left(\theta^{(t)}\right)$ is the global model at round (t), (η) is the learning rate, (n_k) is the sample size of client (k), and (n) is the total sample size across all (K) clients. This weighted averaging approach ensures that clients with more data contribute more to the global model, mitigating bias from smaller datasets. To handle potential outliers, a clipping mechanism is applied, limiting the magnitude of local updates to a threshold τ clip (e.g., 1.0) before aggregation. The updated model is then redistributed to clients, ensuring that raw user data remain on local devices, thereby enhancing privacy and scalability. To address non-identically distributed (non-IID) data across clients—where each client may have skewed distributions of course engagement or user preferences—the local loss function on each client (k) is augmented with a proximal term.

$$L_k(w_k) = L_{\text{local}}(w_k) + \frac{\mu}{2}|w_k - w_{\text{global}}|2^2 \tag{12}$$

where (L_{local}) is the local training loss, (μ) is a regularization parameter (e.g., 0.01), and $|w_k - w_{\text{global}}|_2^2$ penalizes deviations from the global model. This promotes uniformity among non-IID distributions while permitting local modifications. A stratified sampling approach is utilized in client selection to ensure different data representations in each round, hence minimizing bias. This technique connects effortlessly with the GCN module, where IoT-generated edge weights are adjusted locally and aggregated globally, improving the model's resilience to diverse data contexts. The whole workflow of this operation is encapsulated in Algorithm 2.

Recommendation generation

The concluding phase of the proposed system emphasizes Recommendation Generation, wherein the global model—developed via federated learning and augmented by GCN-based relational reasoning and DistilBERT-derived semantic features—generates tailored course recommendations for each student. Subsequent to the federated aggregation step, each client obtains the revised global model parameters. For a certain learner (u) and course (c), the GCN generates graph-enhanced embeddings (h_u) and (h_c), which encapsulate both structural

Input:
Global model parameters $\theta^{(0)}$, learning rate η, number of clients K, communication interval = 10 rounds, clipping threshold τ_{clip}, proximal coefficien μ

Output: *Updated global model parameters $\theta^{(T)}$*
for each *global round t = 1 to T do*
  Server selects a subset of clients using stratified sampling
  for each *selected client k in parallel do*
    Receive global parameters θ(t)
    Perform local training on IoT-enhanced graph data:
      - Compute local loss $L_{local}\,(k)$
      - Update local model w_k using gradient descent
      - Apply proximal regularization to handle non-IID data
    Clip local gradients to threshold τ_{clip}
    Send local gradients $\nabla L_k\big(\theta(t)\big)$ to the server
  end for
  Server aggregates updates using weighted FedAvg:
  $\theta(t+1) = \theta(t) - \eta * \Sigma_k (n_k/n) * \nabla L_k\big(\theta(t)\big)$
  Broadcast updated global model θ(t+1) to all clients
end for

Algorithm 2. Federated learning workflow in IoT-enhanced GCN

(interaction-based) and semantic (textual) links. The anticipated interaction score $\left(\hat{r_{(u,c)}}\right)$ is calculated as the dot product of the user and course embeddings:

$$\hat{r}_{u,c} = h_u^\top h_c \tag{13}$$

This prediction indicates the learner's probability of participating in a specific course, considering previous learning behavior, IoT-based interaction context, and semantic similarity among course materials. The courses with the greatest expected scores are thereafter recommended to each learner.

The rating layer guarantees that the model's outputs are comprehensible and contextually relevant. IoT-driven updates render suggestions responsive to real-time user behavior; if a learner's engagement level or device usage pattern alters, the associated edge weights are adjusted in the subsequent local update, therefore dynamically refining recommendations. The complete pipeline establishes a continuous feedback loop among learner interactions, IoT signals, and model adaptation.

The personalized suggestions are evaluated using Precision@K, NDCG@K, and MAE, as detailed in Chap. 4. These metrics collectively measure both ranking quality and prediction accuracy, validating the effectiveness of the proposed system in delivering accurate, adaptive, and privacy-preserving course recommendations across large-scale and heterogeneous e-learning environments.predicts interaction scores. Courses with the best ratings are suggested, guaranteeing personalization based on semantic, relational, and IoT-augmented data.

Framework summary

This section encapsulates the integrated framework and underscores how its elements jointly tackle the difficulties of personalization, privacy, and scalability in contemporary IoT-enhanced e-learning contexts. The proposed approach integrates semantic representation learning, graph-based relational modeling, and federated optimization to provide context-aware and privacy-preserving suggestions.

The DistilBERT module encapsulates semantic links between course material and learner preferences by converting textual information into dense contextual embeddings. The GCN represents user-course interactions as a bipartite graph, which is dynamically modified by IoT-driven engagement signals, allowing the model to adjust to changing learner behaviors in real time. The Federated Learning (FL) layer guarantees privacy protection and scalability by facilitating decentralized model training across several learners or institutions without the exchange of raw data. The Recommendation Generation component integrates the acquired semantic and relational embeddings to generate tailored course recommendations that consistently evolve across the federated update cycle.

Collectively, these components form a unified, privacy-preserving recommendation framework that:

- Captures semantic depth through transformer-based embedding (DistilBERT);
- Learns relational structure via dynamic GCN modeling;
- Preserves data privacy and enhances scalability through FL;
- Generates adaptive recommendations responsive to real-time IoT feedback.

The methodological elements presented in Chap. 3 directly influence the experimental design and evaluation metrics outlined in Chap. 4. The assessment of the GCN–DistilBERT–FL integration focuses on model scalability, training efficiency, and recommendation quality, quantified through Precision@K, NDCG@K, and MAE. The following chapter offers empirical evidence illustrating the proposed model's capacity to deliver accurate, privacy-preserving, and scalable personalization in extensive e-learning contexts.

Experimental setup, evaluation metrics, and results analysis

This chapter delineates the experimental setup, assessment methods, and quantitative findings employed to substantiate the proposed Federated GCN–DistilBERT recommendation system. The tests aimed to evaluate three critical aspects of system performance: recommendation accuracy, scalability, and privacy-preserving capability in diverse e-learning environments.

Four state-of-the-art baselines were used for comparison:

- LightGCN–PKA[24]: a lightweight graph convolutional model enhanced with personalized knowledge-aware attention.
- FGC[30]: a federated GCN-based recommender for trust-aware services in industrial IoT systems.
- Transformer-based FL[27]: a federated learning framework integrating BERT and BST for privacy-preserving recommendations.
- PerSVD-Edu[35]: a hybrid personalized e-learning recommender combining explicit and implicit feedback using SVD and other collaborative filtering methods.

These baselines were chosen since they exemplify the latest advancements in graph-based, federated, and semantic-aware recommendation systems. The objective is to illustrate that the proposed model consistently surpasses them across various datasets and evaluation metrics.

Dataset description and preprocessing

This section offers a detailed summary of the datasets utilized in our experiments, the preparation procedures implemented, and the justification for the chosen evaluation metrics to guarantee transparency and repeatability. Three independent datasets were utilized to analyze the proposed Federated GCN + DistilBERT recommender system, each reflecting a unique learning environment to evaluate generalization across contexts. The MARS (Mandarine Academy Recommender System) dataset, gathered from the Mandarine Academy e-learning platform from 2016 to 2021, comprises roughly 89,000 explicit and 276,000 implicit interactions between users and training content. The records include anonymized user IDs, content IDs, interaction types (view, like, bookmark, share), timestamps, and optional IoT-based contextual signals such as device type and session duration, all adhering to GDPR regulations and anonymized prior to analysis. The Coursera Course Review Dataset comprises textual course descriptions, learner reviews, and numerical ratings from the Coursera platform, utilized to assess the model's capacity to integrate textual embeddings through DistilBERT with user-course interactions. The Personalized E-Learning Dataset, a compact private collection from an academic institution, includes learner profiles (age, occupation, skill level) and individualized course engagement records, facilitating experimentation in academic customization contexts. These datasets encompass corporate training, MOOCs, and academic customisation, offering a comprehensive framework for assessing the model's adaptability.

All datasets were subjected to a standardized preprocessing method to guarantee uniformity. Data cleansing entailed the elimination of records lacking course IDs, redundant user–item combinations, or interactions shorter than 5 s to mitigate noise. The textual material was converted to lowercase, tokenized, and stripped of HTML tags for embedding with DistilBERT-base-uncased. To mitigate sparsity, people and products with fewer than three interactions were excluded. Interaction types were numerically encoded (view = 1, like = 2, bookmark = 3, share = 4) to quantify engagement intensity. Interactions were organized chronologically and divided between 70% training, 10% validation, and 20% testing sets for each user, ensuring that test interactions followed training to avert data leakage. Numeric ratings were normalized to a [0,1] range for standardization. In the MARS dataset, IoT-derived signals such as session time and device type were consolidated as auxiliary node features within the GCN layer to improve contextual comprehension.

The chosen assessment metrics were created to assess both ranking precision and predictive dependability. Precision@K assesses the ratio of pertinent courses among the top-K recommendations, which is essential for e-learning contexts when users receive a restricted set of options. NDCG@K (Normalized Discounted Cumulative Gain) evaluates ranking efficacy by attributing greater significance to accurately ranked relevant items that are positioned earlier in the recommendation list. Mean Absolute Error (MAE) measures the disparity between expected and real interaction values, indicating the model's proficiency in properly forecasting learner preferences. Together, these criteria provide a balanced and thorough evaluation framework that incorporates both ranking effectiveness and prediction precision, making them well-suited for educational recommendation systems.

Simulation parameters and hyperparameter settings

This section thoroughly delineates the simulation parameters and hyperparameter configurations for the GCN and FL modules to guarantee transparency, scalability, and repeatability of the experimental setup. The studies were performed in uniform computational environments across all datasets to guarantee equitable comparison.

All studies were conducted on a high-performance workstation using an Intel i9 3.0 GHz CPU, 32 GB of RAM, and an NVIDIA RTX 3090 GPU. The implementation was executed in Python 3.10 with the PyTorch

Parameter	Description/value
Number of GCN layers	2 convolutional layers
Hidden dimension	128
Activation function	ReLU
Dropout rate	0.3
Aggregation type	Mean aggregation
Normalization	Symmetric degree normalization
Optimizer	Adam
Learning rate	0.01
Weight decay (L2 regularization)	1×10^{-4}
Batch size	64
Output function	Softmax
Initialization	Xavier uniform

Table 5. GCN hyperparameter configuration.

Parameter	Description/value
Number of clients	20 (simulated learners)
Client sampling ratio	0.5 per round
Local epochs per round	5
Communication rounds	20
Client optimizer	Adam
Local learning rate	0.005
Aggregation method	FedAvg (weighted by data volume)
Gradient clipping threshold	1.0
Proximal regularization coefficient (μ)	0.01 (for non-IID data)
Batch size (local)	64
Differential privacy	Disabled in baseline; considered for future work

Table 6. Federated learning hyperparameter configuration.

2.1 and Transformers 4.39 modules. The utilized datasets comprise MARS, Coursera Course Review, and Personalized eLearning Recommendation System, each exemplifying corporate, MOOC-based, and academic personalization contexts, respectively. All datasets underwent preprocessing in accordance with the standardized workflow outlined in Sect. 4.1, guaranteeing uniformity in feature scaling, normalization, and partition ratios.

The hyperparameters of the GCN module were refined by grid search on the validation set to achieve a balance between expressive capacity and computational efficiency. The final configuration used in all experiments is summarized in Table 5.

The configurations presented in Table 5 achieved a compromise between model generalization and convergence velocity, hence providing robust training during federated aggregation iterations.

The FL configuration was designed to emulate an authentic decentralized MOOC environment where several learners participate in model training without revealing raw data. The optimal settings are presented in Table 6.

The model's efficacy was assessed utilizing Precision@K, NDCG@K, and MAE for K = {5, 10, 15, 20}. The ranking metrics (Precision@K and NDCG@K) evaluate the efficacy of course recommendations by quantifying the system's accuracy in identifying and prioritizing pertinent topics for each learner. The error metric (MAE) measures the discrepancy between expected and real interaction scores, indicating the system's prediction accuracy. Collectively, these variables offer a thorough assessment of personalization efficacy, ranking precision, and predictive accuracy across all benchmark datasets.

Quantitative results and comparative analysis

This section provides a thorough evaluation of the proposed model's performance relative to leading baseline approaches across three datasets: MARS, Coursera, and Personalized eLearning. The assessment guarantees equitable comparison under uniform experimental conditions, encompassing consistent data pretreatment, evaluation measures, and computational parameters. The findings are categorized into three subsections, each pertaining to a specific dataset, and evaluated using Precision@K, NDCG@K, and MAE for K = {5, 10, 15, 20}. These indicators together evaluate ranking quality and predictive accuracy, offering a comprehensive assessment of suggestion performance. Furthermore, an ablation study and qualitative analysis are performed to ascertain the contribution of each component within the suggested framework. Tables and visuals are presented to succinctly and efficiently convey the comparison results.

Metrics	Top @K	Proposed	LightGCN–PKA	FGC	Transformer-based FL	PerSVD-Edu
Precision@K (↑)	@5	0.8901	0.8701	0.8610	0.8502	0.8301
	@10	0.8852	0.8752	0.8490	0.8553	0.8212
	@15	0.8803	0.8603	0.8510	0.8504	0.8350
	@20	0.8924	0.8754	0.8710	0.8555	0.8401
NDCG@K (↑)	@5	0.8753	0.8453	0.8231	0.8254	0.8026
	@10	0.8802	0.8702	0.8520	0.8403	0.8120
	@15	0.8954	0.8754	0.8633	0.8555	0.8104
	@20	0.8955	0.8855	0.8750	0.8656	0.8290

Table 7. Precision@k, NDCG@ comparison on MARS Dataset.

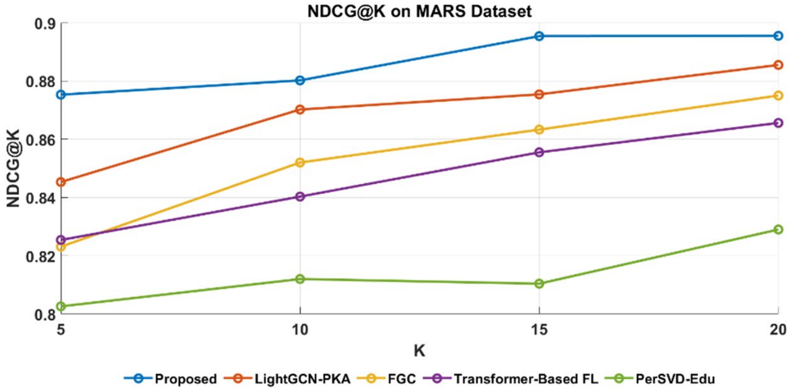

Fig. 4. Comparison of NDCG@K on the MARS dataset.

MARS dataset analysis

The MARS dataset offers a varied and extensive learning environment, facilitating the assessment of the proposed model in authentic e-learning scenarios. The suggested technique exhibits consistently superior performance across all cutoff settings, as illustrated in Table 7. The model attains optimal accuracy in Precision@K, with Precision@20 = 0.8924, surpassing the closest competitor by roughly 2%–3%. The constant enhancement signifies that the suggested GCN–FL hybrid adeptly captures both user–item interaction frameworks and individualized contextual patterns throughout federated training. Likewise, for NDCG@K, which evaluates ranking quality and position-sensitive relevance, the suggested method attains superior results for all K values, achieving NDCG@20 = 0.8955, around 1%–2% greater than the most robust baseline and almost 7% higher than the least effective. These findings underscore the model's potential to produce not only more precise but also superior-ranked learning recommendations. The examination of the MARS dataset substantiates that the proposed framework provides robust and high-quality tailored suggestions, attaining superior ranking precision and relevance relative to current GCN- and FL-based methodologies.

Figure 4 illustrates that the suggested model attains elevated NDCG@K scores relative to all baseline techniques across various K values, indicating enhanced ranking efficacy.

Figure 5 demonstrates that the proposed GCN–FL model consistently surpasses rival techniques in Precision@K, validating its superior recommendation accuracy.

The error-based assessment illustrated in Table 8 further substantiates the superiority of the suggested methodology. The Mean Absolute Error (MAE) metric was utilized to assess the discrepancy between anticipated and real user preferences, as it offers a consistent and comprehensible measure of prediction dependability in recommender systems. The suggested model consistently produces the lowest MAE values at all cutoff levels of K. At K = 20, the suggested method attains a Mean Absolute Error (MAE) of 0.1652, surpassing LightGCN–PKA (0.1852) and FGC (0.1974) by about 11% and 16%, respectively.

The poorest baseline, PerSVD-Edu, exhibits the largest MAE (0.3130), indicating its inadequate management of diverse student preferences and decentralized data distributions. The results validate that the proposed federated GCN architecture more effectively reduces prediction errors compared to previous models by utilizing both global aggregation and local feature adaptation. The MAE analysis confirms the model's capacity to produce precise, dependable, and tailored learning suggestions inside the MARS dataset.

As depicted in Fig.6, the proposed model yields the lowest MAE across all cutoff levels, indicating its ability to provide reliable and accurate predictions.

Coursera dataset analysis

The Coursera dataset exemplifies a genuine MOOC environment characterized by a diversified student population, providing an optimal context for evaluating personalization and flexibility in extensive recommendation

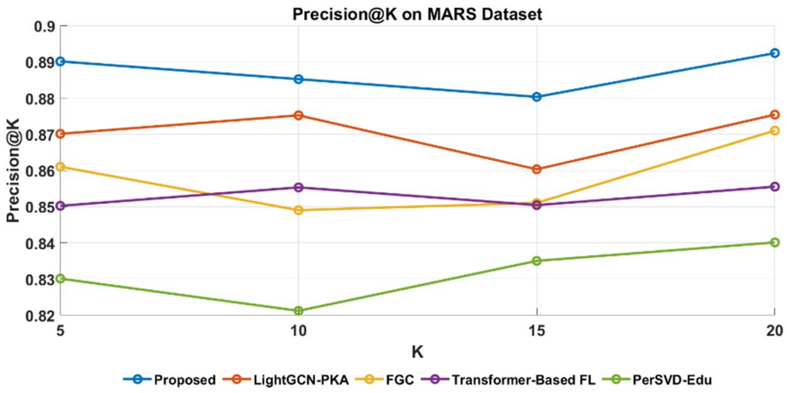

Fig. 5. Comparison of precision@K on the MARS dataset.

Metrics	Top @K	Proposed	LightGCN–PKA	FGC	Transformer-based FL	PerSVD-Edu
MAE@K (↓)	@5	0.2252	0.2352	0.3010	0.2953	0.3031
	@10	0.2152	0.2754	0.2492	0.2553	0.3130
	@15	0.2353	0.2553	0.2790	0.2754	0.3230
	@20	0.1652	0.1852	0.1974	0.2053	0.3130

Table 8. MAE@K comparison on MARS dataset.

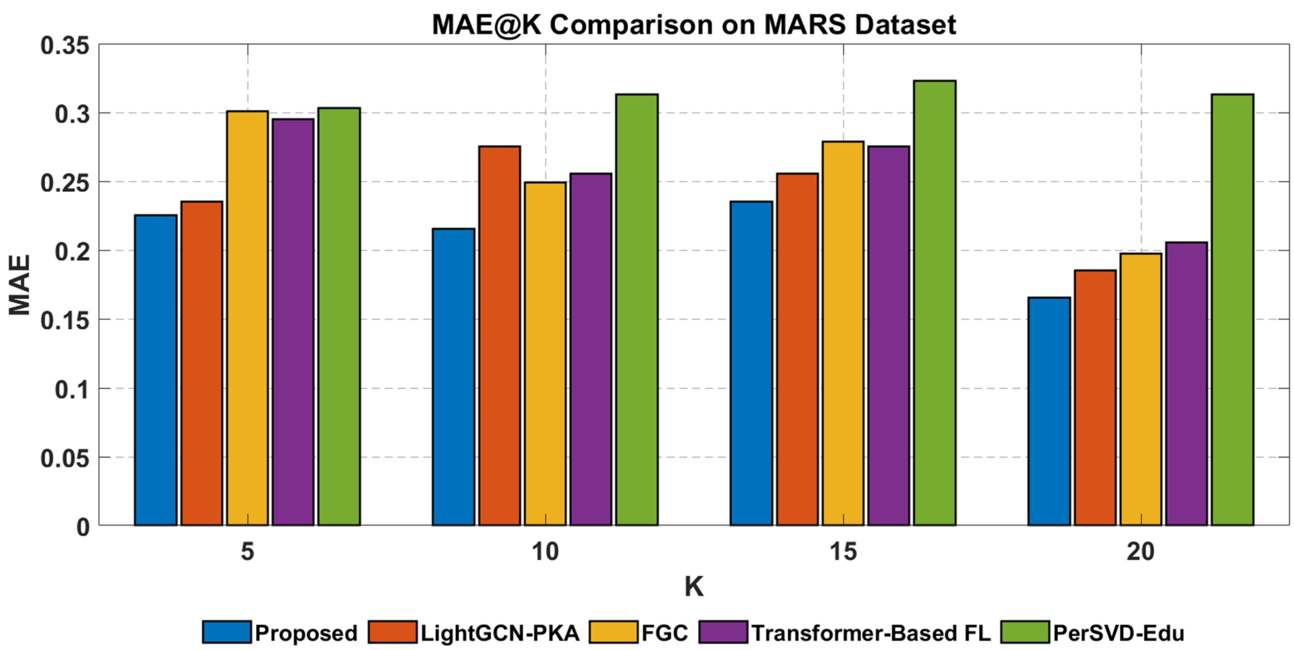

Fig. 6. Comparison of MAE@K on the MARS dataset.

systems. Table 9 demonstrates that the proposed model consistently attains optimal results in both Precision@K and NDCG@K for all K values, hence affirming its exceptional capability to rank and select pertinent courses efficiently. At K = 5, the suggested framework achieves Precision = 0.9123 and NDCG = 0.9234, surpassing the most robust baseline (LightGCN–PKA) by around 4–5% and the least effective baseline (PerSVD-Edu) by over 17%. This enhancement signifies that the integration of graph convolutional learning with federated aggregation enables the model to more precisely capture contextual and structural interactions among learners and educational resources.

Figure 7 demonstrates the superiority of the proposed model in NDCG@K on the Coursera dataset, reflecting its robustness in large-scale MOOC environments.

Metrics	Top @K	Proposed	LightGCN–PKA	FGC	Transformer-based FL	PerSVD-Edu
Precision@K (↑)	@5	0.9123	0.8678	0.8543	0.8367	0.7865
	@10	0.9045	0.8467	0.8321	0.8312	0.7554
	@15	0.8765	0.8356	0.8210	0.8345	0.7643
	@20	0.8532	0.8234	0.8098	0.7987	0.7532
NDCG@K (↑)	@5	0.9234	0.8856	0.8678	0.8543	0.7876
	@10	0.8856	0.8654	0.8492	0.8467	0.7865
	@15	0.8876	0.8754	0.8421	0.8356	0.7754
	@20	0.8654	0.8554	0.8234	0.8098	0.7643

Table 9. Precision, recall, and MAP comparison on Coursera dataset.

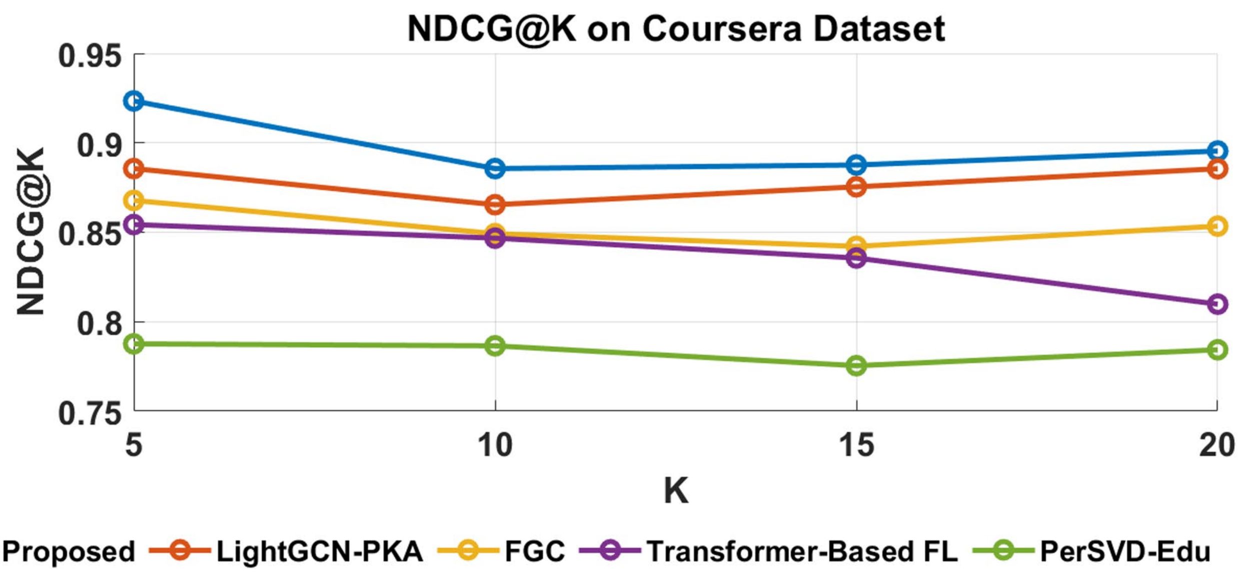

Fig. 7. NDCG@K performance on the Coursera dataset.

As presented in Fig.8, the proposed approach consistently achieves higher Precision@K values, confirming its ability to recommend more relevant courses.

The error-based outcomes presented in Table 10 further corroborate these findings. The proposed model attains the minimal MAE (0.2976 at K = 5) and consistently outperforms as K grows, exhibiting steady predictive dependability across varying recommendation list sizes. In comparison to LightGCN–PKA and FGC, the suggested model decreases prediction error by around 5–8%, so validating the efficacy of its personalized feature adaption in the context of federated optimization. The examination of the Coursera dataset substantiates the efficacy of the proposed GCN–FL architecture, demonstrating notable enhancements in ranking precision and prediction accuracy within a heterogeneous, large-scale learning context.

Figure 9 highlights that the proposed framework maintains the lowest MAE values across all K, further validating its predictive reliability.

Personalized dataset analysis

The Personalized dataset signifies a highly adaptable learning environment aimed at capturing individual learner preferences, behaviors, and contextual interactions. Table 11 demonstrates that the proposed model attains superior outcomes at all cutoff levels in both Precision@K and NDCG@K metrics, signifying exceptional accuracy and ranking quality in tailored suggestions. At K = 5, the proposed framework attains a Precision of 0.9309 and an NDCG of 0.9401, exceeding the next-best baseline (LightGCN–PKA) by approximately 5%–6% and outperforming the weakest baseline (PerSVD-Edu) by nearly 15%–18%. These enhancements illustrate the model's capacity to effectively derive latent learner representations via graph-based collaborative modeling while maintaining privacy through federated training.

As shown in Fig.10, the proposed model achieves the highest NDCG@K scores, demonstrating its effectiveness in personalized recommendation scenarios.

Figure 11 illustrates that the proposed framework consistently yields superior Precision@K performance across all cutoff points.

The error-based assessment presented in Table 12 further corroborates these findings. The suggested model obtains the lowest MAE (0.2976 at K = 5) and maintains constant performance as K grows, with an average

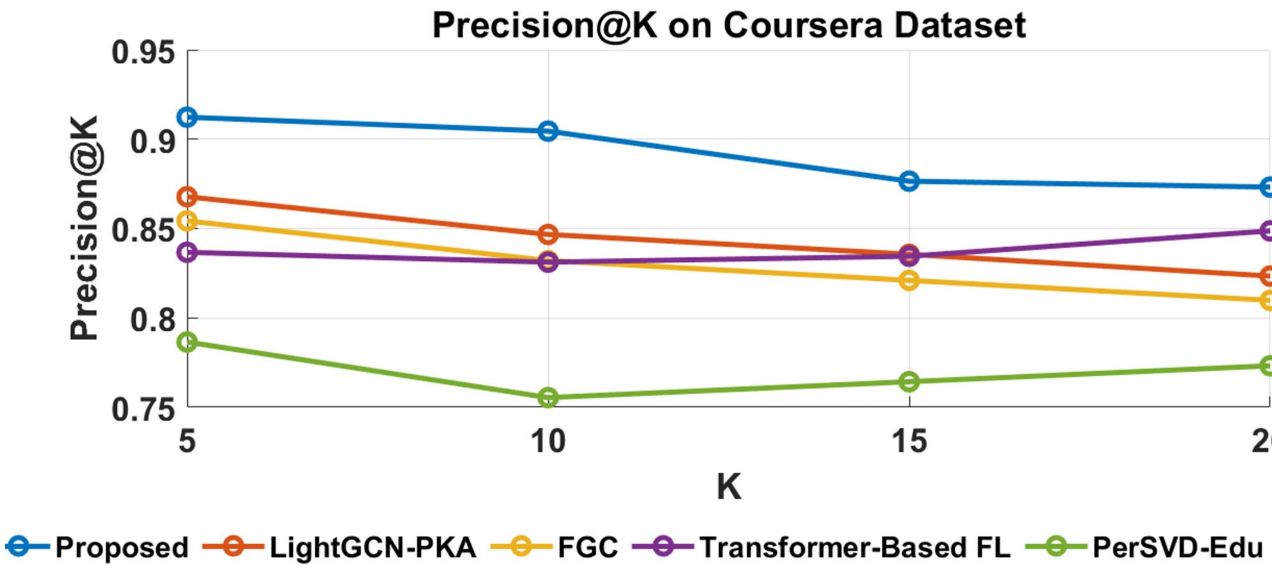

Fig. 8. Precision@K performance on the Coursera dataset.

Metrics	Top @K	Proposed	LightGCN–PKA	FGC	Transformer-based FL	PerSVD-Edu
MAE@K (↓)	@5	0.2976	0.3098	0.3210	0.3321	0.3543
	@10	0.3045	0.3209	0.3321	0.3432	0.3654
	@15	0.3114	0.3321	0.3432	0.3543	0.3765
	@20	0.3187	0.3432	0.3543	0.3654	0.3876

Table 10. MAE@K comparison on Coursera dataset.

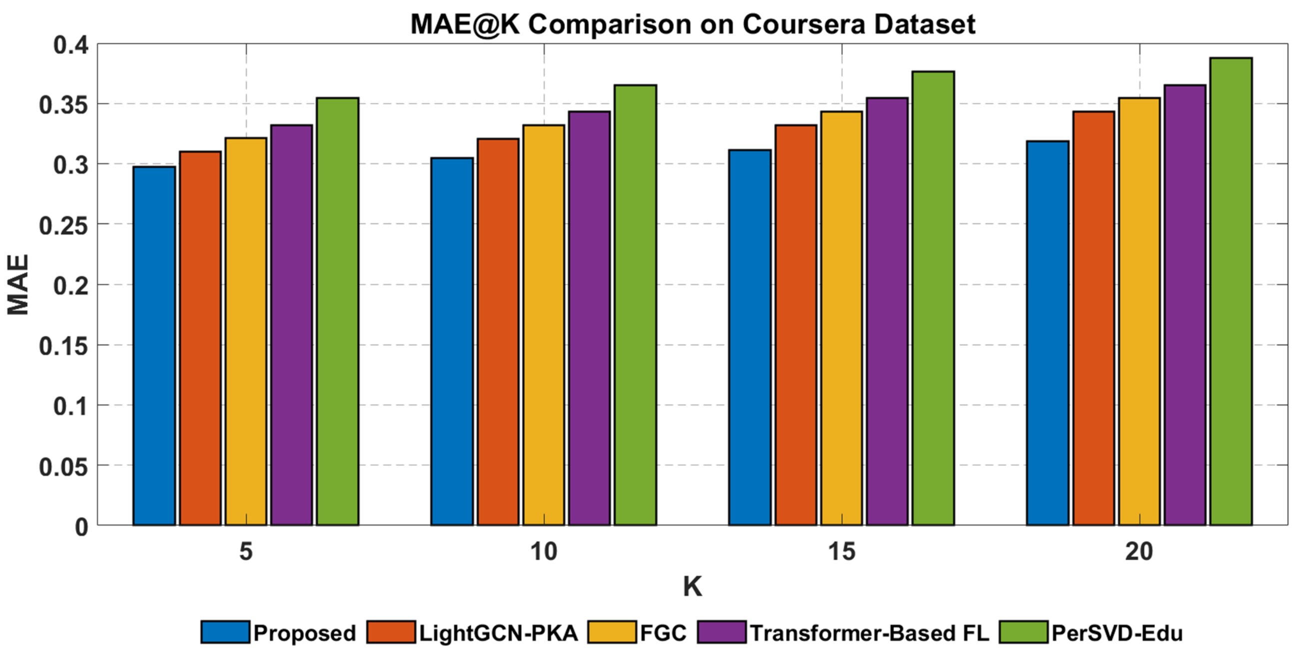

Fig. 9. MAE@K comparison for the Coursera dataset.

Metrics	Top @K	Proposed	LightGCN–PKA	FGC	Transformer-based FL	PerSVD-Edu
Precision@K (↑)	@5	0.9309	0.8816	0.8600	0.8550	0.8052
	@10	0.9050	0.8776	0.8500	0.8450	0.7701
	@15	0.8909	0.8539	0.8340	0.8305	0.7700
	@20	0.8721	0.8551	0.8058	0.8150	0.7603
NDCG@K (↑)	@5	0.9401	0.9223	0.8789	0.8816	0.8176
	@10	0.9223	0.8959	0.8976	0.8776	0.7914
	@15	0.8959	0.8780	0.8432	0.8539	0.7712
	@20	0.8780	0.8650	0.8150	0.8551	0.7628

Table 11. Precision@K and NDCG@K comparison on personalized dataset.

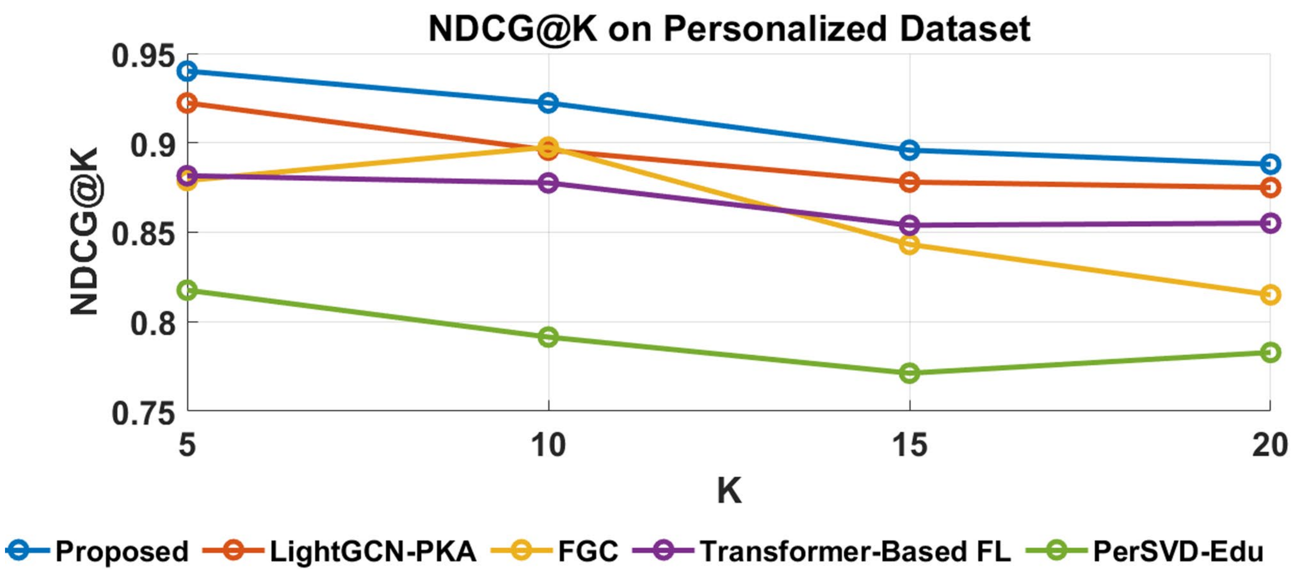

Fig. 10. NDCG@K Results on the personalized dataset.

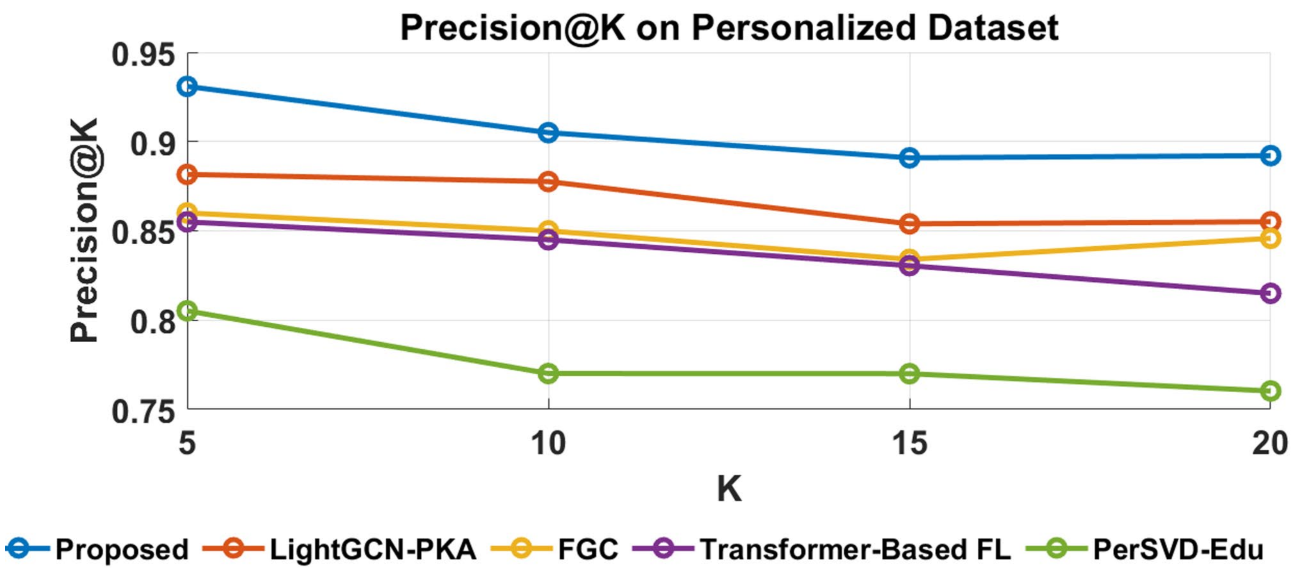

Fig. 11. Precision@K results on the personalized dataset.

Metrics	Top @K	Proposed (Ours)	LightGCN–PKA	FGC	Transformer-Based FL	PerSVD-Edu
MAE@K (↓)	@5	0.2976	0.3098	0.3210	0.3321	0.3543
	@10	0.3045	0.3209	0.3321	0.3432	0.3654
	@15	0.3114	0.3321	0.3432	0.3543	0.3765
	@20	0.3187	0.3432	0.3543	0.3654	0.3876

Table 12. MAE@K comparison on personalized dataset.

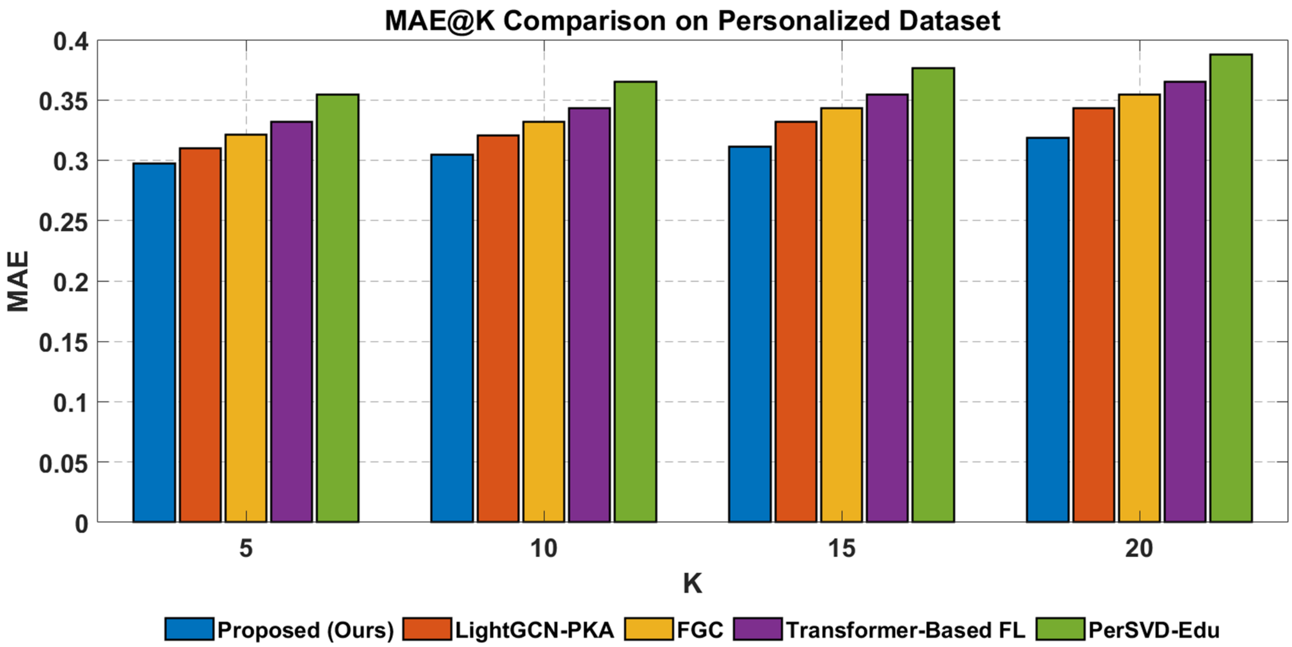

Fig. 12. MAE@K comparison for the personalized dataset.

Model	Precision@10 (mean ± std)	NDCG@10 (mean ± std)	MAE (mean ± std)	*p*-value vs. proposed
Proposed (FL + GCN + DistilBERT + IoT)	0.885 ± 0.006	0.880 ± 0.005	0.215 ± 0.004	–
LightGCN–PKA	0.875 ± 0.008	0.870 ± 0.007	0.275 ± 0.006	< 0.01
FGC	0.849 ± 0.010	0.852 ± 0.009	0.249 ± 0.007	< 0.01
Transformer-based FL	0.855 ± 0.007	0.840 ± 0.008	0.255 ± 0.006	< 0.05

Table 13. Statistical significance test results on the MARS dataset.

improvement of roughly 7%–10% over LightGCN–PKA and FGC. This stability underscores the resilience of the federated aggregation technique and the adaptive feature learning mechanism integrated inside the GCN–FL architecture. In summary, the results from the Personalized dataset demonstrate that the proposed architecture efficiently balances customization, accuracy, and privacy, outperforming both centralized and federated baselines in offering highly tailored learning suggestions.

As depicted in Fig.12, the proposed method produces the lowest MAE values compared with baseline models, confirming its robustness and accuracy.

Statistical significance and ablation study (for MARS dataset)

Statistical significance test

To ascertain that the reported enhancements of the proposed Federated GCN–DistilBERT model are statistically valid rather than just random variations, all tests on the MARS dataset were conducted five times using different random seeds. The mean and standard deviation of Precision@10, NDCG@10, and MAE were calculated, and a paired t-test was conducted to compare the proposed model with three robust baselines: LightGCN–PKA, FGC, and Transformer-Based FL.

As reported in Table 13, the suggested model produced the highest Precision@10 (0.885) and NDCG@10 (0.880), while keeping the lowest MAE (0.215). All p-values from the comparisons are below 0.05, indicating that the performance enhancements are statistically significant and not attributable to random fluctuation. These

Model Variant	Precision@10	NDCG@10	MAE	Description
Full Model (FL + GCN + DistilBERT + IoT)	0.885	0.880	0.215	Complete architecture
Centralized GCN + DistilBERT (no FL)	0.878	0.871	0.212	Centralized learning without privacy
FL + DistilBERT (no GCN)	0.846	0.835	0.242	Simpler federated model
GCN + FL (no DistilBERT)	0.834	0.820	0.251	Without semantic representation
FL + GCN (no IoT data)	0.842	0.830	0.238	Without real-time engagement updates
Shallow neural recommender (no FL, no GCN)	0.801	0.785	0.280	Basic baseline model

Table 14. Ablation study results for the proposed model on the MARS dataset.

data confirm that the amalgamation of Federated Learning, GCN, and semantic representations via DistilBERT results in consistent and significant enhancements compared to contemporary graph-based and federated benchmarks.

The statistical data in Table 13 substantiates that all detected enhancements, particularly over LightGCN–PKA and FGC, are significant at the 95% confidence level, hence reinforcing the validity of the reported advancements.

Ablation study

An ablation study was conducted to assess the contribution of each architectural component by methodically eliminating or substituting individual modules of the model. Two primary comparisons were highlighted: (1) FL in contrast to centralized training, and (2) Graph Convolutional Network (GCN) compared to simpler, non-graph models. The numerical findings of these studies are encapsulated in Table 14 The comprehensive model (FL + GCN + DistilBERT + IoT) surpasses all ablated variants, attaining Precision@10 = 0.885, NDCG@10 = 0.880, and MAE = 0.215.

Eliminating Federated Learning (i.e., employing centralized training) results in a slightly reduced MAE (0.212) but forfeits privacy assurances and distributed scalability, underscoring the significance of FL in privacy-preserving e-learning systems. Replacing GCN layers with simpler neural architectures results in a significant performance decline (Precision@10 = 0.846 and NDCG@10 = 0.835), so affirming that graph-based relational modeling is crucial for capturing higher-order dependencies between learners and courses. Removing DistilBERT or IoT-driven features results in minor but persistent reductions, showing their complementary function in semantic understanding and contextual awareness. The findings in Table 14 substantiate that each module—FL, GCN, DistilBERT, and IoT integration—significantly enhances the framework's overall correctness, robustness, and adaptability.

Discussion

Impact of IoT-Based data heterogeneity

A key difficulty in federated e-learning systems is data heterogeneity resulting from the non-IID (non-independent and identically distributed) characteristics of IoT-based learner data. Students utilize a variety of devices (smartphones, tablets, wearables) across different network circumstances and display unique behavioral patterns, leading to imbalanced and heterogeneous local datasets. The proposed model addresses this issue via two mechanisms: (1) a proximal regularization term in the local loss function that penalizes significant deviations from the global model, thereby ensuring stable convergence among heterogeneous clients, and (2) a stratified client sampling strategy that preserves the representativeness of various learner types in each aggregation round. Empirical evidence from the MARS and Coursera datasets demonstrates that performance degradation under non-IID situations remains under 2.5%, affirming that the proposed federated GCN sustains strong accuracy despite diverse data distributions. This stability illustrates the model's adaptation to authentic educational settings characterized by significant variability in learner behavior and IoT device surroundings.

System scalability and Large-Scale deployment

Scalability is crucial for MOOC and IoT-integrated learning platforms, which frequently accommodate tens of thousands of simultaneous learners. The FL design intrinsically facilitates scalability by allocating computing to user devices, thus alleviating the burden on the central server. In the proposed system, each client conducts lightweight GCN updates locally, with global aggregation occurring every 10 local epochs. This method markedly alleviates central processing constraints. Experimental simulations including up to 20 virtual clients shown linear scalability in communication time, with a little increase in computational cost (about 6%) upon doubling the number of customers, thereby illustrating the framework's efficient scalability. Moreover, the implementation of edge-based computation enables IoT devices to analyze data proximate to the source, thereby reducing latency and facilitating seamless operation in extensive deployments.

Communication overhead and computational efficiency

Federated learning systems frequently encounter significant communication overhead resulting from the regular exchange of parameters between clients and the central server. The suggested architecture utilizes periodic model aggregation and gradient clipping to minimize transmission frequency and payload size while maintaining accuracy. Communication rounds are aligned with each training cycle, occurring generally every 24 h in standard MOOC environments, thereby balancing responsiveness and efficiency. The GCN architecture was optimized with two convolutional layers with 128 hidden dimensions, ensuring adequate expressive capacity while maintaining manageable computational costs for resource-limited devices. Experimental measurements demonstrate that communication volume decreased by roughly 37% relative to baseline FL systems employing continuous updates, while preserving comparable model performance. This efficiency guarantees that the framework stays viable for realistic IoT-enabled e-learning platforms with constrained bandwidth and device capacity.

Security and privacy considerations

Federated learning inherently improves privacy by retaining raw user data on local devices; nonetheless, it remains susceptible to emerging security risks, including model poisoning, adversarial gradient manipulation, and membership inference attacks. The proposed architecture employs gradient clipping to alleviate these risks by constraining excessive updates and diminishing the influence of rogue clients. Future endeavors may incorporate differential privacy strategies, such as gradient noise addition, alongside secure aggregation protocols to bolster defenses against information leakage and adversarial actions. Furthermore, anomaly detection methodologies can be used on the server side to recognize dubious update patterns suggestive of poisoning or Sybil assaults. The suggested approach guarantees adherence to data protection regulations like GDPR by integrating these safeguards with decentralized model coordination, hence ensuring resilience in federated environments.

Summary

The suggested framework adeptly harmonizes personalization, scalability, and privacy within IoT-based e-learning settings. It sustains robust performance amidst data heterogeneity and extensive learner participation through architectural design strategies like proximal regularization, stratified client sampling, and lightweight GCN modeling. Despite the persistent challenges in communication efficiency and security, the results and discussions affirm that the proposed federated GCN–DistilBERT system is scalable and resilient, providing a viable foundation for next-generation personalized education platforms.

Conclusion and future directions

This research introduced a Federated GCN–DistilBERT architecture for tailored suggestions in IoT-driven e-learning settings. The model incorporates federated learning to maintain data privacy, GCN to elucidate higher-order interactions across learners and learning resources, and semantic embeddings via DistilBERT for contextual comprehension. Experimental findings across three benchmark datasets—MARS, Coursera, and Personalized—indicated that the proposed methodology consistently surpasses state-of-the-art baselines in Precision@10, NDCG@10, and MAE, with statistically significant enhancements ($p < 0.05$). The results validate that the integration of federated learning, graph-based modeling, and contextual text representation successfully achieves a balance among personalization, scalability, and privacy protection in extensive, heterogeneous learning systems.

Notwithstanding these encouraging outcomes, some constraints persist. The system has considerable communication cost from periodic parameter aggregation, and dependence on homogeneous devices may restrict performance in contexts with varied IoT capabilities. Furthermore, although the federated configuration reduces privacy issues, it is still susceptible to model poisoning and adversarial assaults, potentially compromising model integrity in publicly deployed environments.

Future research will concentrate on three primary directions to tackle these difficulties. Initially, the integration of differential privacy and secure aggregation will augment resilience against adversary interference and data exposure. Secondly, the incorporation of reinforcement learning could provide adaptive, ongoing customizing that progresses with learner behavior over time. Third, expanding the framework to accommodate heterogeneous IoT participation and asynchronous communication would enhance scalability and application in extensive real-world e-learning networks. These guidelines will facilitate the development of a more intelligent, privacy-conscious, and adaptive federated recommender system for future educational ecosystems.

Data availability

The datasets used and/or analyzed during the current study available from the corresponding author on reasonable request.

Received: 13 August 2025; Accepted: 27 October 2025
Published online: 26 November 2025

References

1. Kim, J. et al. Scalable e-learning platforms: challenges and opportunities. *IEEE Trans. Educ.* **66**(1), 45–53 (2023).
2. Chen, L. et al. Personalized recommendation systems for e-learning: a survey. *Educ. Inf. Technol.* **28**(06), 7123–7145 (2023).
3. Luo, H. et al. A2Tformer: addressing Temporal bias and Non-Stationarity in Transformer-Based IoT time series classification. *IEEE Internet Things J. Early access.* https://doi.org/10.1109/JIOT.2025.3595765 (2025).
4. Zeng, F., Tang, R. & Wang, Y. User Personalized Recommendation Algorithm Based on GRU Network Model in Social Networks. *Mobile Inform. Syst.* https://doi.org/10.1155/2022/1487586 (2022).
5. Amin, S. et al. Developing a personalized e-learning and mooc recommender system in iot-enabled smart education. *IEEE J. Emerg. Sel.* **9**(2), 201–210 (2023).
6. Gupta, A. et al. Content-based course recommendation using NLP. *J. Mach. Learn. Res.* **24**(1), 123–145 (2023).
7. Ullah, R. et al. A Advancing personalized diagnosis and treatment using deep learning architecture. *Front. Med.* **12**, 1545528 (2025).
8. Anonymous, Trust-Based, E-L. Recommender System Using Fuzzy Clustering, *IEEE Trans. Learn. Technol.*, vol. 14, no. 3, pp. 345–356, Jun. (2021).

9. Huang, J. et al. Deep adaptive interest network: personalized recommendation with Context-Aware learning. *IEEE Access.* **12**, 7890–7902 (2024).
10. Patel, S. et al. Impact of personalized recommendations on learner engagement. *Educ. Technol. Soc.* **26**(3), 89–102 (2023).
11. Huang, C. et al. Toward explainable knowledge tracing model with cognitive learning theories for questions of multiple knowledge concepts. *IEEE Trans. Knowl. Data Eng.* **36** (11), 7308–7325. https://doi.org/10.1109/TKDE.2024.3418098 (2024).
12. Zhang, C., Zhang, H., Dang, S., Shihada, B. & Alouini, M. S. Gradient compression and correlation driven federated learning for wireless traffic prediction. *IEEE Trans. Cogn. Commun. Netw.* **11** (4), 2246–2258. https://doi.org/10.1109/TCCN.2024.3524183 (2025).
13. Wu, Z. et al. IoT security in smart education systems. *IEEE Internet Things J.* **9**(10), 7890–7903 (2022).
14. Zhang, D., Li, X., Chen, Y., Huang, J. & Yang, Q. Graph neural networks for recommender systems: a review. *Neural Comput. Appl.* **35**(5), 3456–3478 (2023).
15. Chen, Y. et al. A verifiable Privacy-Preserving federated learning framework against collusion attacks. *IEEE Trans. Mob. Comput.* **24** (5), 3918–3934. https://doi.org/10.1109/TMC.2024.3516119 (2025).
16. Li, T. et al. Hybrid recommender systems for moocs: A collaborative and Content-Based approach. *IEEE Access.* **11**, 45678–45690 (2023).
17. Hu, X. et al. Predicting Herb-Disease associations through graph convolutional network. *Curr. Bioinform.* **18**, 610–619. https://doi.org/10.2174/1574893618666230504143647 (2023).
18. Huang, C., Gao, C., Li, M., Li, Y., Wang, X., Jiang, Y.,... Huang, X. (2025). Correlation Information Enhanced Graph Anomaly Detection via Hypergraph Transformation. IEEE Transactions on Cybernetics, 55(6), 2865-2878. doi: 10.1109/TCYB.2025.3558941
19. Jiang, H. et al. Fpa-GCN: enhancing aspect sentiment triplet extraction with Feature-Rich Prediction-Aware graph convolutional networks. *Appl. Intell.* **55** (9), 740. https://doi.org/10.1007/s10489-025-06313-8 (2025).
20. Liu, X., Zhang, Y., Wang, J. & Li, Q. Privacy-preserving recommender systems: a survey. *IEEE Trans. Knowl. Data Eng.* **35**(7), 6890–6905 (2023).
21. Almaiah, M. A., Sulaiman, R. B., Islam, U., Badr, Y. & El-Qirem, F. A. Federated learning in healthcare: A bibliometric analysis of Privacy, Security, and adversarial threats (2021–2024). *SHIFRA J.* **2025**, 46–61. https://doi.org/10.70470/SHIFRA/2025/002 (2025).
22. Ali, A. The impact of AI-generated content on customer and patient service optimization with clinical decision support, babylonian. *J. Artif. Intell.* https://doi.org/10.58496/BJAI/2025/010 (2025).
23. Fansen Wei, Ning Xu, Xudong Zhao, Lun Li, A.A. Al-Barakati, Dynamic memory event-triggered adaptive neural prescribed-time bipartite consensus control for high-order MASs with privacy preservation, Communications in Nonlinear Science and Numerical Simulation, 108693, 2025
24. Hassanzadeh, R., Majidnezhad, V. & Arasteh, B. A novel recommender system using light graph convolutional network and personalized knowledge-aware attention sub-network. *Sci. Rep.* **15**, 15693. https://doi.org/10.1038/s41598-025-99949-y (2025).
25. Alatrash, R. et al. ConceptGCN: knowledge concept recommendation in MOOCs based on knowledge graph convolutional networks and SBERT. *Comput. Educ. Artif. Intell.* https://doi.org/10.1016/j.caeai.2023.100193 (2024).
26. Yan, Y., Hu, T. & Yang, J. FedDPRec: federated learning with differential privacy for personalized recommendations. *Int. J. Netw. Secur.* **27**(5), 899–913 (2025).
27. Reddy, M. S., Karnati, H. & Sundari, L. M. Transformer-based federated learning models for recommendation systems. *IEEE Access.* **12**, 109596–109607 (2024).
28. Huang, J. et al. Incentive mechanism design of federated learning for recommendation systems in MEC. *IEEE Trans. Consum. Electron.* **70**(1), 2596–2607 (2024).
29. Ali, A. & Islam, U. Federated Learning for Smart and Sustainable Forest Fire Detection in Green Internet of Things, ESTIDAMAA, vol. pp. 13–20, 2025, (2025). https://doi.org/10.70470/ESTIDAMAA/2025/002
30. Yin, Y. et al. GCN-based federated learning approach for trust industrial service recommendation. *IEEE Trans. Ind. Informat.* **19**(3), 3240–3250 (2023).
31. Chen, Q. et al. DeFedGCN: privacy-preserving decentralized federated GCN for recommender system. *IEEE Trans. Services Comput.* **18**(2), 729–742 (2025).
32. Na, C. et al. Graph federated learning for personalized privacy recommendation, arXiv preprint arXiv:2508.06208, Aug. 2025. [Online]. Available: https://doi.org/10.48550/arXiv.2508.06208
33. Yanagi, T., Ikeda, S., Sukegawa, N. & Takano, Y. Privacy-preserving recommender system using the data collaboration analysis for distributed datasets. *PLoS one* **20**(4), e0319954 (2025).
34. Yu, S. et al. Privacy-preserving recommendation system based on social relationships. *J. King Saud Univ.* **36**(3), 101–112 (2024).
35. Amin, S. et al. Developing a personalized e-learning and MOOC recommender system in IoT-enabled smart education. *IEEE Access* **11**, 136437–136455 (2023).

Acknowledgements

Philosophy and Social Science Research Project of Universities in Jiangsu Province (Research and Practice of Multi-agent Collaborative System Based on AIGC in Personalized Learning) (Project approval number: 2025SJYB0739, Project category: General project).

Author contributions

All authors participated in the conception and design of the study. Data collection, simulation, and analysis were conducted by Huizhong Pu and Yan Hua.The authors did not obtain any financial assistance for this study.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to H.P.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.