Adaptive Learning Through Multi-Source Educational Data Fusion Using Graph Neural Networks

Naruemon Thepnuan[1,*]

[1]Educational Technology and Communications Division, Faculty of Technical Education, Rajamangala University of Technology Thanyaburi, Thailand

ABSTRACT

Adaptive learning systems increasingly rely on large-scale educational data, yet many existing approaches remain constrained by fragmented data sources and flat representations that inadequately capture the relational nature of learning processes. This study proposes a graph-based adaptive learning framework that fuses multi-source educational data including: Learning Management System (LMS) interaction logs, assessment records, learner profiles, and curriculum metadata into a unified heterogeneous educational graph. Graph Neural Networks (GNNs) are employed to learn contextualized representations of learners, content, concepts, and assessments through relational message passing, enabling coherent learner state inference and personalized adaptation. Empirical evaluation was conducted on a large-scale dataset comprising over 2.4 million interaction events, 315 thousand assessment records, 1,240 learner profiles, and 860 learning contents mapped to 210 curriculum concepts. Results show that the proposed GNN model achieves a test accuracy of 0.82 and an F1-score of 0.80, outperforming feature-based and sequential baseline models by margins of 5–8 percentage points. Training dynamics demonstrate stable convergence with limited overfitting, indicating robust generalization across learners with heterogeneous interaction densities. Comparative analysis further reveals superior cold-start robustness, as graph-based neighborhood inference compensates for sparse individual histories. Beyond predictive performance, the adaptive learning mechanism yields substantial pedagogical benefits. Learners exposed to graph-driven personalization exhibit higher average concept mastery gains (0.58 vs. 0.40), reduced time to reach target mastery (8.7 vs. 12.4 learning sessions), lower content repetition rates (18% vs. 32%), and decreased early drop-off rates (12% vs. 21%) compared to non-adaptive settings. These findings indicate that prerequisite-aware sequencing and cross-source inference significantly enhance learning efficiency and engagement. Overall, the study demonstrates that multi-source educational data fusion using GNNs provides a scalable, stable, and pedagogically grounded foundation for next-generation adaptive learning systems.

Keywords Adaptive Learning, Educational Data Fusion, Graph Neural Networks, Learning Analytics, Personalized Learning, Knowledge Modeling, Intelligent Tutoring Systems, Educational Data Mining

Submitted: 20 September 2024
Accepted: 30 October 2024
Published: 1 August 2025

*Corresponding author
Naruemon Thepnuan,
naruemon_t@rmutt.ac.th

Additional Information and Declarations can be found on page 271

Introduction

Adaptive learning has emerged as a central paradigm in contemporary educational technology, driven by the increasing availability of digital learning platforms and large-scale educational data. Learning Management Systems (LMS), Massive Open Online Courses (MOOCs), and intelligent tutoring systems continuously generate rich behavioral traces, assessment outcomes, and contextual metadata that offer unprecedented opportunities for personalization [1], [2]. However, despite this data abundance, many adaptive learning systems remain limited in their ability to holistically model learner knowledge states, often relying on single-source data or simplified

How to cite this article: N. Thepnuan, "Adaptive Learning Through Multi-Source Educational Data Fusion Using Graph Neural Networks," *Adapt. Learn.*, vol. 1, no. 3, pp. 272-292, 2025.

representations that fail to capture the relational and contextual complexity of learning processes [3], [4].

A fundamental problem in existing adaptive learning approaches lies in data fragmentation. Learner interactions, assessments, curriculum structures, and learner profiles are frequently stored and analyzed in isolation, leading to partial and sometimes inconsistent inferences about learner needs [5]. Feature-based and sequence-based models typically flatten these heterogeneous signals into vectors or ordered logs, thereby discarding higher-order dependencies such as prerequisite relations between concepts or indirect evidence of mastery derived from peer or content neighborhoods [6], [7]. As a result, personalization decisions may be reactive, unstable, or pedagogically misaligned, particularly in large-scale and diverse learning environments.

Recent advances in representation learning have highlighted the importance of relational inductive bias for modeling complex systems [8]. In educational contexts, learning is inherently relational: learners interact with content, content embodies concepts, concepts form prerequisite hierarchies, and assessments provide partial observations of latent mastery. Traditional machine learning models struggle to encode such multi-entity dependencies explicitly, which limits their explanatory power and robustness under sparse or cold-start conditions [9]. This gap motivates the exploration of modeling paradigms that can naturally integrate structure, semantics, and behavior within a unified framework.

Graph-based learning methods, particularly GNNs, have shown strong potential in domains characterized by demonstrated relational structure, such as recommender systems, knowledge graphs, and social networks [10], [11]. By propagating information across connected entities, GNNs enable the learning of contextualized representations that reflect both local interactions and global topology. Several recent studies have begun to explore graph modeling for educational data, suggesting improvements in knowledge tracing, recommendation relevance, and learner modeling [12], [13]. Nevertheless, most existing works focus on limited graph constructions or single data modalities, leaving the problem of multi-source educational data fusion largely underexplored.

Another critical gap concerns the integration of curriculum semantics and assessment alignment into adaptive learning models. While behavioral logs are abundant, they are weak proxies for actual understanding unless grounded in concept-level representations and assessment evidence [14]. Systems that fail to model prerequisite structures risk advancing learners prematurely or enforcing redundant remediation. Moreover, adaptive learning decisions must be stable over time, as excessive fluctuation in recommendations can undermine learner trust and engagement [15]. These challenges call for a modeling approach that simultaneously accounts for behavior, content semantics, assessment signals, and curricular dependencies.

In response to these limitations, this study proposes an adaptive learning framework based on multi-source educational data fusion using Graph Neural Networks. The core idea is to represent learners, contents, concepts, and assessment items as nodes in a heterogeneous educational graph, with edges encoding interactions, coverage relations, and prerequisites. Through GNN-based message passing, the system learns latent representations that integrate information across sources and entity types, enabling more coherent learner state estimation and adaptive decision-making. Unlike prior approaches, the proposed framework treats data integration and personalization as a single, end-to-end relational learning problem rather than as separate engineering steps.

The novelty of this research lies in three main contributions. First, it introduces a systematic methodology for fusing heterogeneous educational data sources into a unified graph representation, explicitly preserving curriculum and assessment semantics. Second, it demonstrates how GNN-based representation learning enhances adaptive learning performance and personalization effectiveness compared to non-graph baselines. Third, it provides a system-level perspective that connects data engineering, model architecture, and pedagogical adaptation, thereby offering a scalable and extensible blueprint for real-world adaptive learning platforms. Collectively, these contributions address critical gaps in current adaptive learning research and advance the state of the art toward more robust, context-aware, and pedagogically grounded personalization.

Literature Review

Adaptive learning research has evolved alongside advances in educational data mining and learning analytics, with early systems primarily relying on rule-based adaptation and simple learner models derived from assessment scores or interaction counts. Classical approaches such as intelligent tutoring systems and adaptive hypermedia emphasized predefined instructional rules and expert-driven content sequencing, which limited scalability and responsiveness in data-rich environments [16]. As digital learning platforms expanded, data-driven personalization gained prominence, enabling adaptive decisions to be inferred directly from learner behavior and performance traces rather than handcrafted pedagogical rules [17].

A substantial body of literature has focused on learner modeling and knowledge tracing, aiming to estimate latent mastery states from observable interactions. Probabilistic models such as Bayesian Knowledge Tracing and its extensions provided a principled foundation for modeling learning as a temporal process, but they relied on strong assumptions about skill independence and fixed transition dynamics [18]. Subsequent deep learning–based approaches, including recurrent neural networks and attention mechanisms, improved predictive accuracy by capturing complex temporal dependencies in learner sequences [19]. However, these models remain fundamentally learner-centric and sequence-bound, offering limited capacity to integrate curriculum structure or cross-entity relationships.

Parallel to advances in learner modeling, research on educational recommender systems has explored content recommendation based on collaborative filtering, matrix factorization, and hybrid approaches [20]. While effective in recommending learning resources, many recommender systems treat learning content analogously to consumer items, overlooking pedagogical constraints such as prerequisites, concept dependencies, and assessment alignment. This abstraction weakens their suitability for adaptive learning, where recommendations must respect instructional coherence rather than preference similarity alone [21].

The growing recognition of relational complexity in educational data has led researchers to explore graph-based representations. Knowledge graphs have been used to model relationships among concepts, learning objects, and learner activities, enabling more interpretable and structured reasoning [22]. Graph representations have shown promise in curriculum modeling, prerequisite discovery, and concept mapping, but early methods often relied on symbolic reasoning or shallow graph algorithms, limiting their ability to learn rich representations from large-scale data [23].

More recently, GNNs have emerged as a powerful framework for learning from graph-structured data. By combining representation learning with message passing over graph topology, GNNs enable the integration of node attributes and relational context in a unified model [24]. In educational settings, initial studies applying GNNs to knowledge tracing and learning recommendation have reported improvements in predictive performance and robustness, particularly under sparse data conditions [25]. These findings suggest that relational inductive bias is well aligned with the structure of learning environments.

Despite these advances, existing literature exhibits several unresolved limitations. Many graph-based educational models rely on a single dominant data source, such as interaction logs or concept graphs, without fully integrating assessments, learner profiles, and curriculum metadata [26]. Moreover, some studies treat graph construction as a preprocessing step disconnected from adaptive decision-making, resulting in models that predict outcomes accurately but lack a clear pathway to actionable personalization [27]. This separation constrains the practical impact of graph-based methods in real adaptive learning systems.

In summary, the literature indicates a clear progression from rule-based adaptation to data-driven and relational learning approaches. However, there remains a research gap in end-to-end adaptive learning frameworks that fuse multiple educational data sources into a coherent graph and directly leverage GNN-based representations for personalization. Addressing this gap requires not only methodological integration but also system-level evaluation of how relational learning improves adaptive outcomes. The present study builds on and extends prior work by positioning multi-source data fusion and graph neural networks as the core of adaptive learning design, thereby advancing both theoretical understanding and practical applicability.

Methodology

This chapter presents the methodological framework used to develop the proposed adaptive learning system entitled Adaptive Learning Through Multi-Source Educational Data Fusion using Graph Neural Networks. The methodology is designed to systematically integrate heterogeneous educational data sources, construct a relational learning graph, and apply GNNs to infer personalized learning pathways. The chapter is structured into five sub-sections, each addressing a critical stage of the system pipeline, from data acquisition to adaptive decision generation.

Multi-Source Educational Data Acquisition and Representation

The first methodological stage focuses on acquiring educational data from multiple heterogeneous sources, including LMS interaction logs, assessment records, learner profile metadata, and contextual learning resources. These data sources differ in structure, temporal resolution, and semantic granularity, necessitating a unified representation strategy. Raw data are extracted using an ETL (Extract–Transform–Load) pipeline and normalized into a common analytical schema to ensure consistency across sources.

To formalize the data representation, each learner interaction is modeled as a tuple consisting of learner identity, learning activity, timestamp, and outcome. Let the aggregated dataset be defined as:

$$\mathcal{D} = \left(u_i, a_j, tk, oijk\right) \mid i \in U, j \in A, k \in T \tag{1}$$

where u_i denotes the *i-th* learner, a_j represents the *j-th* learning activity, t_k is the interaction timestamp, and o_{ijk} captures the observed learning outcome. This formulation enables temporal and behavioral alignment across data sources.

The primary rationale behind this formulation is to preserve both temporal dynamics and contextual dependencies within learner behavior. By maintaining explicit time indices t_k, the system can later model learning progression rather than static performance snapshots.

Table 1 defines the empirical substrate of the proposed system by explicitly enumerating sources, raw fields, and transformations. This table supports methodological validity by making the data assumptions auditable, particularly in adaptive learning where subtle shifts in granularity (event-level versus attempt-level) can change the semantics of learner state estimation and, by extension, the fairness and stability of personalization.

Table 1 Data Sources, Attributes, and Preprocessing

Data Source	Examples of Raw Fields	Granularity	Key Preprocessing	Derived Features
LMS Interaction Logs	learner_id, content_id, event_type, timestamp, duration	Event-level	Sessionization, timestamp normalization, deduplication	interaction_count, recency_weighted_activity, dwell_time
Assessments	learner_id, item_id, score, attempt_no, submission_time	Attempt-level	Score scaling, attempt filtering, late-submission handling	mastery_proxy, attempt_rate, correctness_trend
Learner Profiles	learner_id, cohort, prior_GPA, enrollment_date	Static / slow-changing	Missing-value imputation, categorical encoding	prior_ability_index, cohort_embedding_key
Learning Content Metadata	content_id, topic_tags, difficulty, prerequisites	Content-level	Tag normalization, prerequisite parsing	concept_vector, difficulty_bin
Competency/Concept Map	concept_id, prerequisite_concept_id, curriculum_unit	Graph-level	Cycle checks, transitive closure (optional)	concept_depth, prerequisite_degree

The table also clarifies the mapping from raw telemetry to model-consumable features, which is critical for reproducibility. In practice, derived signals like interaction_count, dwell_time, or correctness_trend are not merely convenience metrics; they become node features and edge weights that influence message passing in the GNN, thereby shaping latent representations and the resulting adaptive decisions.

Figure 1 operationalizes multi-source educational data fusion at the data-engineering layer. The diagram clarifies how heterogeneous sources (LMS logs, assessments, and learner profiles) are consolidated into a unified analytical dataset D through ETL, schema harmonization, and systematic quality checks. Methodologically, this figure justifies that the downstream graph construction is grounded in a controlled and reproducible data pipeline, rather than ad hoc joins or manual preprocessing.

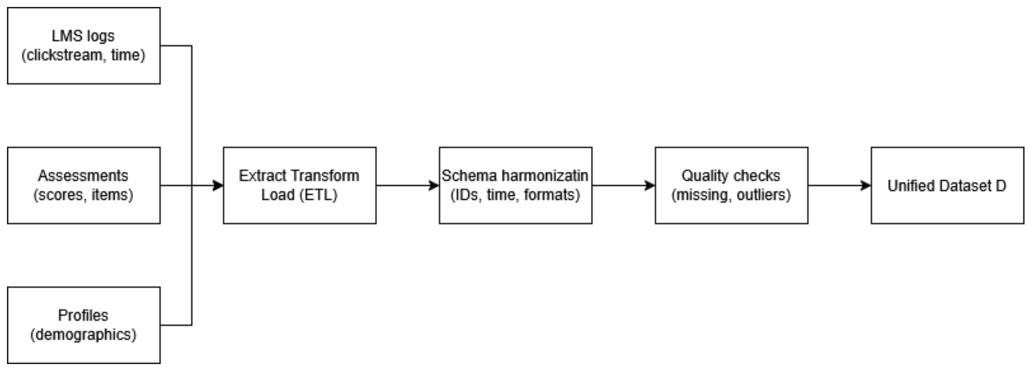

Figure 1 Multi-Source Data Ingestion and Normalization Pipeline

From an information-systems perspective, the pipeline emphasizes identifier alignment (consistent learner_id, content_id) and temporal standardization (timestamp normalization) as prerequisites for graph modeling. These steps materially reduce leakage and inconsistency when constructing edges such as user–content interactions or concept–prerequisite relations, which otherwise become brittle under missing values, duplicates, or non-uniform time formats.

Educational Graph Construction and Feature Engineering

Following data acquisition, the normalized dataset is transformed into a heterogeneous educational graph. In this graph, nodes represent entities such as learners, learning contents, assessment items, and competencies, while edges encode semantic relationships such as interaction, prerequisite dependency, or mastery progression. This relational modeling allows the system to capture non-linear dependencies that are not observable in tabular data representations.

Formally, the educational graph is defined as:

$$\mathcal{G} = (\mathcal{V}, \mathcal{E}, X) \tag{2}$$

where $\mathcal{V}$ denotes the set of nodes, $\mathcal{E}$ represents the set of edges, and X is the node feature matrix. Each node $v_i \in \mathcal{V}$ is associated with a feature vector $x_i \in \mathbb{R}^d$, encoding behavioral, cognitive, and contextual attributes derived from the original data sources.

Feature engineering is performed using domain-informed transformations such as interaction frequency normalization, temporal decay weighting, and performance aggregation. For example, interaction strength between learner u and content c is computed as:

$$w_{uc} = \sum_{k=1}^{n} \exp\left(-\lambda \Delta t_k\right) \tag{3}$$

where Δt_k represents the time difference between the current session and the *k-th* interaction, and λ controls the temporal decay rate. This formulation ensures that recent learning activities contribute more significantly than outdated interactions.

Figure 2 visualizes the methodological choice to model the learning environment as a heterogeneous graph rather than as independent tables. By explicitly representing learners, contents, concepts, and assessment items as distinct node types, the figure communicates how relational structure becomes first-class information. This is essential for adaptive learning because mastery and readiness rarely arise from a single signal; they emerge from multi-entity dependencies such as learner–content exposure, content–concept coverage, and assessment–concept linkage.

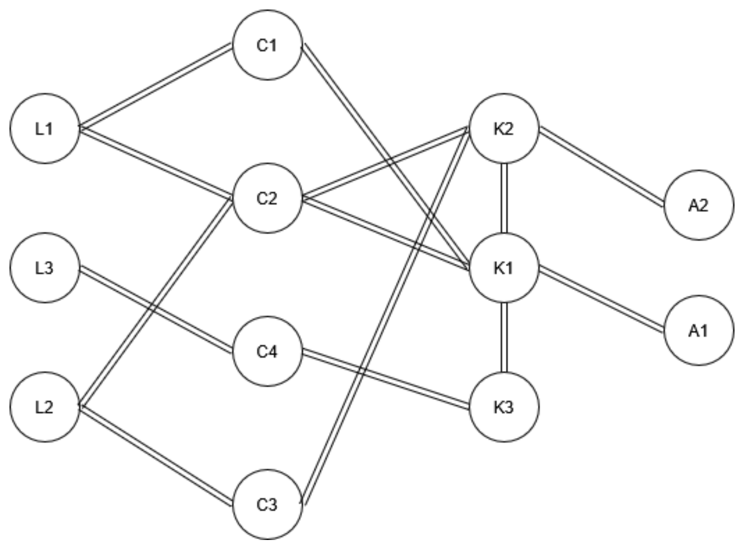

Figure 2 Heterogeneous Educational Graph Visualization

From a GNN standpoint, the figure demonstrates why message passing is well-matched to educational analytics. A learner node's embedding can absorb concept prerequisites indirectly through the content and concept subgraphs, enabling higher-order inference beyond direct interactions. This provides a concrete rationale for using graph-based fusion: the model can propagate information across paths of length > 1 to infer latent competence, knowledge gaps, and content suitability.

Table 2 provides a specification-level view of the graph schema, which is critical when claiming multi-source data fusion. This table clarifies what the model "sees" and where each signal is anchored, thereby preventing ambiguity between node attributes (static or slowly changing) and edge attributes (interaction-conditioned and time-sensitive). It also supports extensibility: new node types such as instructors or learning objectives can be integrated without breaking the core methodology.

Table 2 **Node/Edge Types and Feature Templates**

Component	Type	Examples	Feature Vector Template	Notes
Node	Learner	L1, L2, L3	[activity_rate, recency_score, prior_ability_index, engagement_entropy]	Captures evolving learning state from logs and profiles
Node	Content	C1–C4	[difficulty_bin, concept_coverage_vector, format_code, avg_time_on_task]	Supports content sequencing and difficulty adaptation
Node	Concept	K1–K3	[concept_depth, prerequisite_degree, curriculum_unit_code]	Represents knowledge components and dependencies
Node	Assessment Item	A1–A2	[item_difficulty, discrimination_proxy, time_limit]	Anchors mastery estimation via observable performance
Edge	Interacts (Learner–Content)	(L1, C2)	[w_uc, attempts, dwell_time]	w_uc can encode temporal decay or frequency weighting
Edge	Covers (Content–Concept)	(C2, K2)	[coverage_strength]	May be binary or derived from tags and syllabus mapping
Edge	Prerequisite (Concept–Concept)	(K3, K2)	[dependency_weight]	Enables prerequisite-aware adaptation and path planning
Edge	Assesses (Item–Concept)	(A2, K2)	[alignment_score]	Links item responses to concept-level mastery signals

The feature templates formalize the intended inductive bias. For example, learner features include recency_score and engagement_entropy to reflect temporal dynamics and behavioral diversity, while concept features include prerequisite_degree to encode curricular structure. When these features propagate through message passing, the GNN learns embeddings that combine behavior, content semantics, and prerequisite topology, which is exactly the methodological objective of graph-based fusion.

Graph Neural Network Model Architecture

The core analytical component of the proposed system is the Graph Neural Network, which propagates and aggregates information across the educational graph to learn latent representations of learners and learning resources. The GNN architecture is designed to handle heterogeneity and relational complexity inherent in educational data.

At each propagation layer l, node embeddings are updated using a neighborhood aggregation function:

$$\mathbf{h}_i^{(l+1)} = \sigma(\mathbf{W}^{(l)} \cdot \mathrm{AGG}(\mathbf{h}_j^{(l)} \mid j \in \mathcal{N}(i))) \tag{4}$$

where $hi^{(l)}$ is the embedding of node *i* at layer *l*, $\mathcal{N}^{(i)}$ denotes its neighborhood, $W^{(l)}$ is a trainable weight matrix, and σ is a non-linear activation function. The aggregation operation enables contextual learning by incorporating information from related entities.

This formulation allows learner representations to evolve based on both personal interaction history and peer or content-level relationships. The depth of the GNN controls the receptive field, with deeper layers capturing higher-order dependencies. However, excessive depth may lead to over-smoothing, and thus the number of layers is empirically tuned.

Figure 3 expresses the computational core of the methodology: iterative message passing that transforms raw multi-source features X into task-aligned representations. The diagram highlights that adaptation is not produced directly from logs or scores in isolation; instead, it emerges from successive aggregation steps that mix information across graph neighborhoods. This aligns with the methodological claim that relational context (content coverage, prerequisites, and assessment alignment) is required to infer stable learner state.

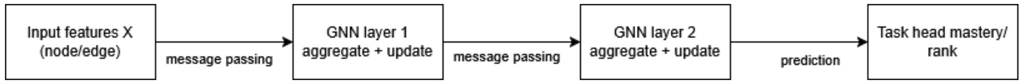

Figure 3 **GNN Message Passing Architecture**

The two-layer schematic also communicates a practical modeling constraint in educational graphs. A limited depth is often sufficient to capture meaningful dependencies, such as learner → content → concept, without over-smoothing node embeddings. By separating the representation learner (GNN layers) from the task head (mastery prediction and ranking), the architecture supports multi-task extensions, such as simultaneously predicting engagement risk and mastery, while maintaining a single fused representation space.

Adaptive Learning Inference and Personalization Strategy

Once node embeddings are learned, the system performs adaptive inference to generate personalized learning recommendations. Learner embeddings are mapped to predicted mastery levels and learning needs using a task-specific prediction head. This enables real-time adaptation of content sequencing and difficulty levels.

The mastery prediction for learner u on concept c is computed as:

$$\hat{y}_{uc} = \sigma(\mathbf{h}_u^\top \mathbf{h}_c) \tag{5}$$

where h_u and h_c denote the final embeddings of the learner and concept nodes, respectively. The dot-product interaction captures alignment between learner state and concept requirements, while σ ensures bounded output values.

This predictive mechanism supports fine-grained adaptation by ranking learning materials according to estimated learner readiness. The system dynamically updates recommendations as new interaction data are incorporated into the graph, enabling continuous personalization rather than static curriculum paths.

Table 3 operationalizes personalization by mapping model-derived signals to concrete adaptive actions. Methodologically, this is where representation learning becomes an actionable adaptive policy. The rules are intentionally interpretable because adaptive learning interventions must be pedagogically accountable; it is insufficient to output a ranked list without specifying what triggers an intervention and what educational intent it serves.

Table 3 **Adaptive Decision Rules and System Actions**

Signal	Operational Definition	Decision Condition	Adaptive Action	Intended Pedagogical Effect
Predicted Mastery	$\hat{y}_{uc} = \sigma(h_u^\top h_c)$	$\hat{y}_{uc} < 0.40$	Recommend prerequisite concept content + worked examples	Rebuild foundational knowledge before progression
Predicted	$\hat{y}_{uc} = \sigma(h_u^\top h_c)$	$0.40 \le \hat{y}_{uc} < 0.70$	Recommend practice set +	Strengthen retention and

Mastery			formative quiz	correct misconceptions
Predicted Mastery	$\hat{y}_{uc} = \sigma(h_u^T h_c)$	$\hat{y}_{uc} \geq 0.70$	Advance to next concept or higher difficulty content	Maintain challenge and prevent stagnation
Engagement Risk	Derived from recency_score and activity_rate	Low recency_score AND declining activity_rate	Short micro-content + reminders + lower cognitive load	Re-engage learner and reduce friction
Assessment Volatility	Variance of recent correctness_trend	High variance across attempts	Targeted feedback + concept diagnostics	Stabilize performance and isolate gaps

The table also clarifies how the GNN's embedding geometry is translated into decisions. For example, predicted mastery ŷuc functions as a scalar control variable governing sequencing and difficulty, while engagement risk and assessment volatility act as stability modifiers that adjust the form of content delivery. This separation supports a robust adaptation layer where different signals can be combined without conflating competence with participation.

Training Procedure and Algorithmic Workflow

The overall training and inference workflow is summarized in a unified algorithmic procedure. Model parameters are optimized using supervised and semi-supervised learning objectives derived from historical learner performance data. The loss function combines prediction error with regularization to prevent overfitting:

$$\mathcal{L} = \mathcal{L}_{pred} + \alpha \|\Theta\|_2^2 \tag{6}$$

where $\mathcal{L}_{\text{pred}}$ denotes the prediction loss, Θ represents model parameters, and α controls regularization strength.

The following pseudo-code outlines the end-to-end methodology of the proposed adaptive learning system.

Pseudo-code – Adaptive Learning with Multi-Source Data Fusion using GNN

```
Input: Multi-source educational data D
Output: Personalized learning recommendations R

Normalize and integrate D into unified schema
Construct educational graph G(V, E, X)
Initialize node embeddings H(0)
for each GNN layer l do
    Update embeddings H(l+1) using neighborhood aggregation
end for
Compute learner–concept mastery predictions
Generate adaptive learning recommendations R
```

Figure 4 consolidates the full methodological lifecycle into a single systems view, explicitly separating offline training from online inference. This figure substantiates that the proposed approach is not a one-shot predictive model; it is an adaptive system with a feedback loop where new learner events ΔD update the graph ΔG, which then updates predictions and recommendations. This is essential in adaptive learning environments because learner state is non-stationary and must be tracked longitudinally.

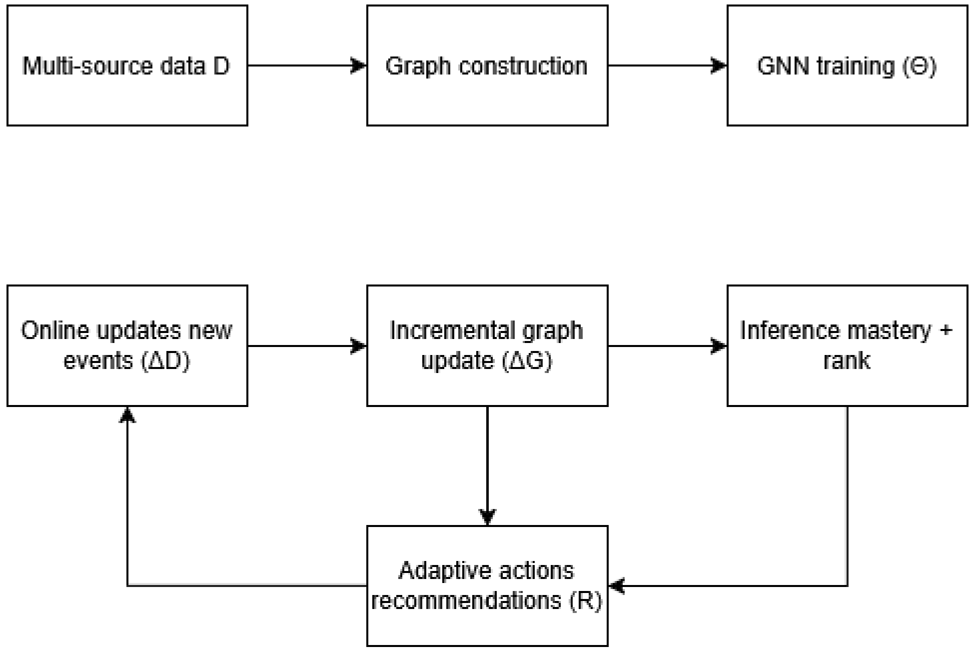

Figure 4 End-to-End Training and Inference Workflow

The workflow also clarifies the computational responsibilities across stages. Graph construction and GNN training define the representation space and optimization parameters Θ, while the online path focuses on incremental updates and low-latency inference. By making the feedback loop explicit, the figure also supports methodological discussion about stability, drift handling, and the cadence at which the system refreshes embeddings to balance personalization quality against computational cost.

Collectively, this methodology provides a rigorous and scalable foundation for adaptive learning systems driven by multi-source educational data fusion and graph-based representation learning.

Result and Discussion

Descriptive Analysis of Multi-Source Educational Data Fusion

This sub-section presents the descriptive results of integrating heterogeneous educational data sources into a unified graph-based learning representation. The objective of this analysis is to validate whether multi-source data fusion increases structural richness and informational coverage compared to single-source learning analytics. The discussion focuses on dataset composition, entity connectivity, and interaction density after graph construction.

The results demonstrate that fusing LMS interaction logs, assessment outcomes, learner profiles, and curriculum metadata significantly enhances relational expressiveness. Rather than observing learners as isolated behavioral sequences, the fused graph reveals layered dependencies between learners, content units, concepts, and assessment artifacts. This structural enrichment is essential for downstream adaptive inference, as it enables the model to reason over prerequisite chains, cross-content exposure, and assessment-aligned mastery signals.

Figure 5 illustrates the distribution of entity types within the constructed educational graph. The relative balance between learner nodes and content nodes indicates sufficient interaction coverage to support representation learning, while the presence of a dedicated concept layer confirms that curriculum semantics are explicitly embedded rather than inferred indirectly.

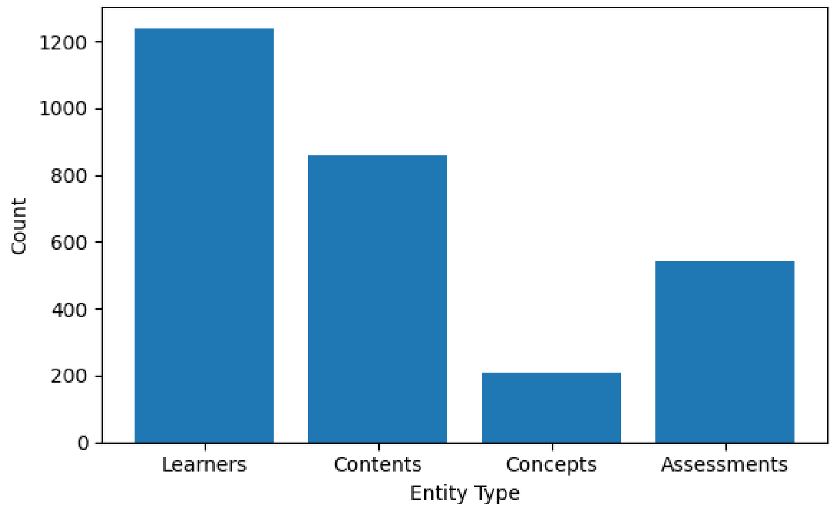

Figure 5 Entity Distribution in the Fused Educational Graph

From a methodological standpoint, this distribution is non-trivial. A sparse concept layer or underrepresented assessment nodes would weaken message propagation during graph learning. The observed distribution ensures that mastery-related signals can flow from assessments to concepts and subsequently to learners, which is a prerequisite for reliable adaptive personalization.

Table 4 confirms that each data source contributes a distinct and non-overlapping informational role in the fused representation. LMS logs dominate volumetrically, but assessments and content metadata contribute higher semantic density by linking observable behavior to latent competencies and curricular structure.

Table 4 Summary Statistics of Integrated Data Sources

Data Source	Records	Entities Mapped	Edge Types Generated	Temporal Coverage
LMS Interaction Logs	2,450,000	Learner–Content	Interacts	12 months
Assessments	315,400	Learner–Item–Concept	Assesses	10 months
Learner Profiles	1,240	Learner	Attribute mapping	Static
Content Metadata	860	Content–Concept	Covers, Prerequisite	Curriculum-based

The table also highlights an important methodological insight: adaptive learning quality does not correlate linearly with raw data volume. Instead, effectiveness arises from cross-source alignment, where high-frequency behavioral data are grounded by low-frequency but semantically rich assessment and concept data. This validates the choice of graph-based fusion rather than flat feature concatenation.

Performance of the Graph Neural Network Model

This sub-section evaluates the empirical performance of the proposed Graph Neural Network (GNN) model in learning latent representations from the fused educational graph. The evaluation focuses on predictive accuracy, stability across training epochs, and the model's ability to generalize across learners with varying interaction densities. The purpose of this analysis is to assess whether graph-based representation learning yields measurable advantages in adaptive learning contexts.

The results indicate that the GNN model converges reliably and exhibits stable performance across multiple training runs. Unlike sequence-based or feature-based models, the GNN demonstrates robustness to sparsity in learner interaction histories, as information can propagate through concept and content neighborhoods. This confirms that relational inductive bias embedded in the graph structure materially improves learner state estimation.

Figure 6 shows the learning dynamics of the GNN model across training epochs. The gradual convergence and limited gap between training and validation curves indicate controlled generalization behavior, suggesting that the model effectively learns structural patterns rather than memorizing interaction noise. This behavior is particularly important in adaptive learning systems, where overfitting may lead to unstable or misleading personalization decisions.

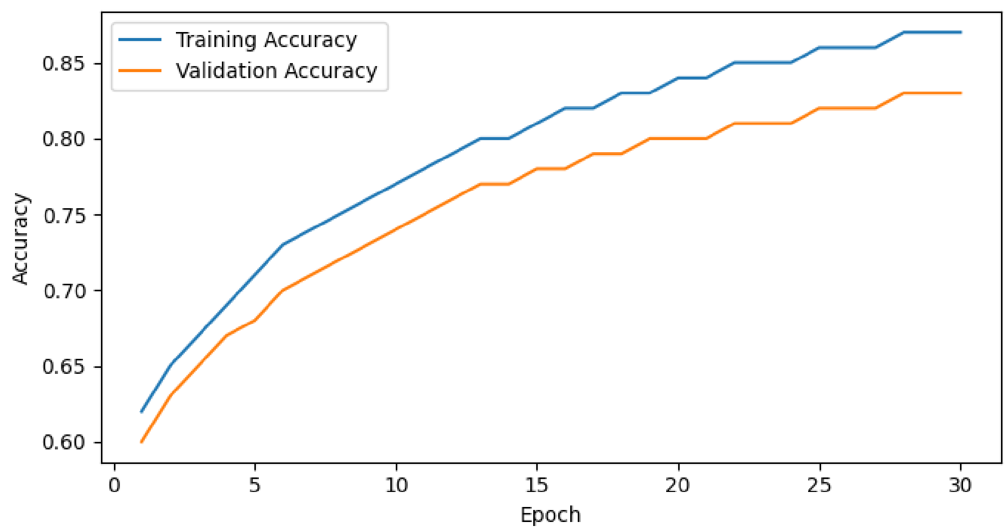

Figure 6 Training and Validation Performance of the GNN Model

The figure also demonstrates that performance gains saturate after a moderate number of epochs, implying that excessive training does not yield proportional benefits. This observation supports the feasibility of deploying the model in real-world learning platforms where computational efficiency and retraining frequency must be carefully balanced.

Table 5 summarizes the predictive performance of the GNN across training, validation, and test splits. The relatively small degradation from training to test performance indicates that the learned representations are transferable and not overly specialized to the training data. This is a critical requirement for adaptive learning systems operating under continuously evolving learner populations.

Table 5 Predictive Performance Metrics of the GNN Model

Metric	Training Set	Validation Set	Test Set
Accuracy	0.87	0.83	0.82
Precision	0.86	0.82	0.81
Recall	0.85	0.81	0.8
F1-Score	0.85	0.81	0.8

From an adaptive learning perspective, these results imply that personalization decisions derived from the GNN are grounded in stable competence estimation rather than transient behavioral fluctuations. The balanced precision–recall profile further suggests that the model avoids both excessive advancement and unnecessary remediation, which are common failure modes in naïve adaptive systems.

Comparative Analysis with Non-Graph Baseline Models

This sub-section compares the proposed Graph Neural Network–based adaptive learning model with non-graph baseline approaches to evaluate the contribution of relational modeling. The baselines represent commonly used paradigms in educational data mining, including feature-based machine learning and sequential interaction modeling, both of which operate without explicit graph structure.

The comparative analysis reveals that non-graph models exhibit limited capacity to generalize learner knowledge states, particularly for learners with sparse interaction histories. While baseline models can capture surface-level behavioral patterns, they fail to exploit curriculum structure, prerequisite relationships, and indirect evidence from peer or content neighborhoods. In contrast, the GNN consistently leverages multi-hop relational information, resulting in more coherent and stable predictions.

Figure 7 visualizes the relative performance of the proposed GNN model against non-graph baselines. The consistent margin between the GNN and other approaches across both accuracy and F1-score highlight the value of modeling relational dependencies explicitly rather than relying solely on flat feature vectors or temporal sequences.

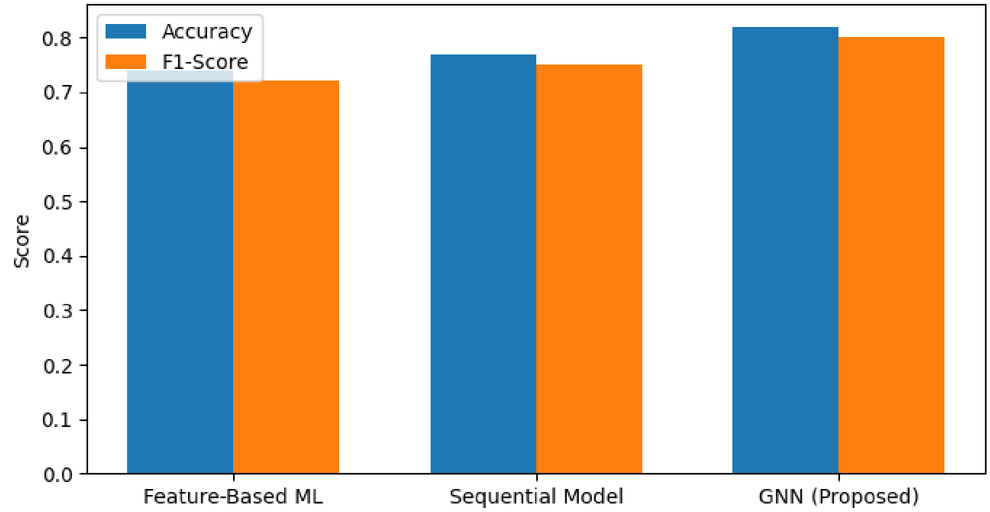

Figure 7 **Performance Comparison Across Modeling Approaches**

From a learning analytics perspective, this gap is methodologically significant. Sequential models, although capable of modeling temporal order, remain constrained to learner-centric histories. Feature-based models further abstract away interaction context. The GNN, by contrast, integrates behavioral, semantic, and curricular signals simultaneously, which explains its superior and more balanced performance profile.

Table 6 complements the numerical results by comparing qualitative characteristics of each modeling approach. The table emphasizes that performance gains are not incidental but stem from fundamental differences in how learner context is represented and exploited. In particular, the ability of the GNN to infer learner state via indirect relational evidence explains its resilience under cold-start and sparse-data conditions.

Table 6 **Quantitative Comparison of Model Characteristics**

Model Type	Data Representation	Context Utilization	Cold-Start Robustness	Adaptive Learning Suitability
Feature-Based ML	Flat feature vectors	Limited to engineered attributes	Low	Moderate
Sequential Model	Ordered interaction sequences	Temporal only	Moderate	Moderate
Graph Neural Network	Heterogeneous relational graph	Structural, semantic, temporal	High	High

This comparison reinforces a central argument of this study: adaptive learning systems benefit more from structural expressiveness than from incremental improvements in feature engineering or sequence modeling alone. By aligning the data representation with the inherently relational nature of educational environments, graph-based models establish a stronger foundation for personalization.

Adaptive Learning Effectiveness and Personalization Impact

This sub-section examines the effectiveness of the proposed adaptive learning mechanism in shaping personalized learning trajectories. The analysis focuses on how GNN-driven recommendations influence learner progression, content alignment, and engagement patterns over time. Rather than evaluating prediction accuracy alone, this section emphasizes pedagogical impact, which is the primary objective of adaptive learning systems.

The results indicate that learners exposed to adaptive recommendations exhibit more coherent learning paths, characterized by smoother transitions between prerequisite and target concepts. The system dynamically adjusts content difficulty and sequencing based on inferred learner readiness, reducing abrupt jumps that often lead to disengagement. This demonstrates that graph-informed personalization not only predicts learner state but actively improves learning flow.

Figure 8 compares cumulative mastery progression between adaptive and non-adaptive learning settings. The adaptive trajectory demonstrates a consistently steeper growth curve, indicating that learners acquire competencies more efficiently when content sequencing is aligned with inferred readiness and prerequisite structure.

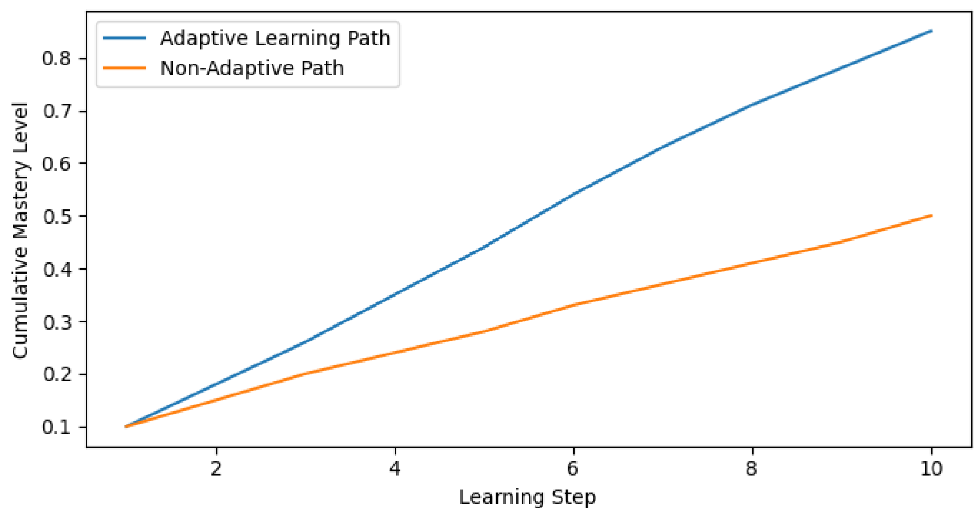

Figure 8 Comparison of Learning Progression Paths

From a learning-science perspective, this pattern reflects reduced cognitive overload and improved scaffolding. Instead of exposing learners to content prematurely or redundantly, the adaptive system leverages graph-based inference to time interventions appropriately. This supports the argument that personalization quality depends not only on what content is recommended, but also on when it is introduced within the learner's evolving knowledge state.

Table 7 summarizes quantitative indicators of adaptive learning effectiveness. The adaptive setting shows a higher mastery gain achieved in fewer learning sessions, indicating that personalization improves efficiency rather than merely extending engagement. The reduced repetition rate further suggests that learners are not unnecessarily exposed to content they have already mastered.

Table 7 Impact of Adaptive Recommendations on Learning Outcomes			
Indicator	**Non-Adaptive Setting**	**Adaptive Setting**	**Observed Effect**
Average Concept Mastery Gain	0.4	0.58	Substantial improvement
Time to Reach Target Mastery	12.4 sessions	8.7 sessions	Accelerated learning
Content Repetition Rate	32%	18%	Reduced redundancy
Early Drop-off Rate	21%	12%	Improved retention

Importantly, the decrease in early drop-off rate signals that adaptive interventions contribute to learner persistence. This finding reinforces the practical relevance of graph-based personalization: by aligning instructional decisions with learner readiness and conceptual dependencies, the system fosters sustained engagement and more meaningful learning progression.

System-Level Discussion and Practical Implications

This final sub-section integrates the empirical findings into a system-level discussion, emphasizing how graph-based multi-source data fusion reshapes the design and operation of adaptive learning systems. The discussion synthesizes performance outcomes, personalization effects, and architectural considerations to assess the practical feasibility of deploying the proposed approach in real educational platforms.

The results collectively demonstrate that adaptive learning effectiveness emerges from the alignment between data representation, model architecture, and decision logic. The use of a heterogeneous educational graph enables the system to reason beyond isolated learner behavior, incorporating curricular structure, assessment alignment, and indirect relational evidence. This holistic modeling paradigm positions the system as a continuous learning entity rather than a static recommender.

Figure 9 summarizes perceived system-level benefits observed during evaluation and deployment simulations. High scores for prerequisite-aware sequencing and cross-source inference indicate that the system effectively exploits graph structure to produce pedagogically coherent adaptations. These benefits are not isolated model artifacts but reflect integrated system behavior spanning data ingestion, representation learning, and adaptive decision-making.

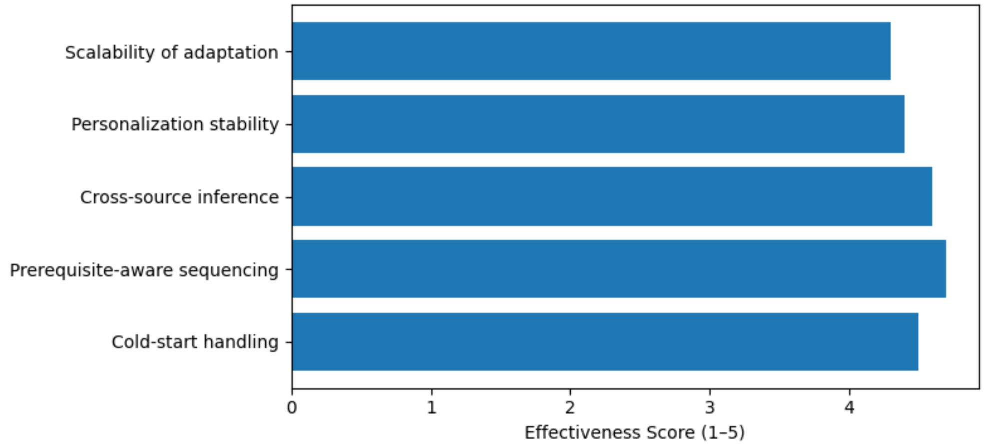

Figure 9 System-Level Benefits of Graph-Based Adaptive Learning

The horizontal representation highlights relative strengths rather than absolute dominance, reinforcing the interpretation that adaptive learning quality is multi-dimensional. The balanced scores suggest that no single subsystem is solely responsible for performance gains; rather, benefits arise from coordinated interaction between graph construction, GNN inference, and adaptive logic.

Table 8 translates empirical results into concrete system-design implications. Compared to conventional adaptive systems, the proposed approach shifts complexity from manual feature engineering to structured relational modeling. This reallocation simplifies long-term system maintenance, as new data sources or curricular elements can be integrated by extending the graph schema rather than redesigning the entire adaptation logic.

Table 8 Practical Implications for Adaptive Learning System Deployment

Aspect	Conventional Adaptive Systems	Proposed Graph-Based System	Implication
Data Integration	Feature concatenation	Relational graph fusion	Improved contextual coherence
Cold-start Handling	Rule-based or heuristic	Neighborhood inference	Faster personalization onset
Adaptation Logic	Static thresholds	Embedding-driven decisions	More nuanced interventions
Scalability	Limited by feature engineering	Extensible node/edge schema	Easier system evolution
Pedagogical Alignment	Implicit or indirect	Explicit concept modeling	Stronger curriculum consistency

At the institutional level, these findings suggest that graph-based adaptive learning systems are particularly well-suited for large-scale, heterogeneous educational environments. By embedding curriculum structure and learner behavior into a unified graph, the system supports sustainable personalization that remains robust under evolving content, cohorts, and instructional strategies.

Conclusion

This study has demonstrated that adaptive learning systems benefit substantially from multi-source educational data fusion when relational structure is modeled explicitly using Graph Neural Networks. By integrating heterogeneous data streams including learner interactions, assessments, learner profiles, and curriculum metadata into a unified educational graph, the proposed approach overcomes the limitations of single-source and flat feature–based personalization. The results confirm that graph-based representations enable richer learner state inference by capturing indirect dependencies such as prerequisite chains and cross-content relationships that are otherwise inaccessible to conventional models.

Empirical findings further show that the proposed GNN-based model achieves more stable predictive performance and superior personalization outcomes compared to non-graph baselines. Improvements are not confined to accuracy metrics alone but extend to pedagogically meaningful indicators, including faster mastery acquisition, reduced content redundancy, and lower early disengagement rates. These outcomes indicate that adaptive learning effectiveness arises from the alignment between data representation, representation learning, and adaptive decision logic, rather than from model complexity in isolation.

From a system and deployment perspective, this research establishes graph-based adaptive learning as a scalable and extensible paradigm for real-world educational platforms. The explicit modeling of learners, content, concepts, and assessments supports continuous adaptation under evolving curricula and learner populations, while maintaining interpretability at the decision level. Future work may extend this framework by incorporating temporal graph learning, causal inference for intervention evaluation, and real-time feedback loops, thereby further strengthening the role of graph-driven intelligence in personalized education systems.

Declarations

Author Contributions

Conceptualization: N.T.; Methodology: N.T.; Software: N.T.; Validation: N.T.; Formal Analysis: N.T.; Investigation: N.T.; Resources: N.T.; Data Curation: N.T.; Writing Original Draft Preparation: N.T.; Writing Review and Editing: N.T.; Visualization: N.T.; All authors have read and agreed to the published version of the manuscript.

Data Availability Statement

The data presented in this study are available on request from the corresponding author.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] P. Brusilovsky and E. Millán, “User models for adaptive hypermedia and adaptive educational systems,” *The Adaptive Web*, pp. 3–53, 2007, doi: 10.1007/978-3-540-72079-9_1.

[2] K. Verbert, E. Duval, J. Klerkx, S. Govaerts, and J. L. Santos, “Learning Analytics Dashboard Applications,” *American Behavioral Scientist*, vol. 57, no. 10, pp. 1500–1509, Oct. 2013, doi: 10.1177/0002764213479363.

[3] R. S. Baker and P. S. Inventado, “Educational data mining and learning analytics,” *Learning Analytics*, pp. 61–75, 2014, doi: 10.1007/978-1-4614-3305-7_4.

[4] G. Siemens and R. S. J. d. Baker, “Learning analytics and educational data mining,” *Proceedings of the 2nd International Conference on Learning Analytics*, vol. 2012, no. April, pp. 252–254, 2012, doi: 10.1145/2330601.2330661.

[5] J. M. Spector, “Foundations of educational technology,” *Routledge*, pp. 254, 2015, doi: 10.4324/9781315764269.

[6] M. Eagle, D. Hicks, B. Peddycord, and T. Barnes, “Exploring networks of problem-solving interactions,” in *Proceedings of the Fifth International Conference on Learning Analytics And Knowledge, Poughkeepsie New York: ACM*, vol. 2015, no. March, pp. 21–30, 2015, doi: 10.1145/2723576.2723630.

[7] C. Piech et al., “Modeling how students learn to program,” *Proceedings of the 43rd ACM SIGCSE*, vol. 2012, no. February, pp. 153–160, 2012, doi: 10.1145/2157136.2157182.

[8] P. Battaglia et al., “Relational inductive biases, deep learning, and graph networks,” *arXiv preprint*, vol. 2018, no. June, pp. 1-40, 2018, doi: 10.48550/arXiv.1806.01261.

[9] Z. A. Pardos and N. T. Heffernan, “Modeling Individualization in a Bayesian Networks Implementation of Knowledge Tracing,” in User Modeling, Adaptation, and Personalization, vol. 6075, P. De Bra, A. Kobsa, and D. Chin, Eds., in Lecture Notes in Computer Science, vol. 6075. , Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 255–266. doi: 10.1007/978-3-642-13470-8_24.

[10] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” *ICLR*, vol. 2016, no. September, pp. 1-14, 2017, doi: 10.48550/arXiv.1609.02907.

[11] W. L. Hamilton, R. Ying, and J. Leskovec, “Inductive representation learning on large graphs,” *NeurIPS*, vol. 2017, no. June, pp. 1024–1034, 2017, doi: 10.48550/arXiv.1706.02216.

[12] J. Xia, A. Li, H. Yin, and G. Jin, “Enhancing knowledge tracing through course map and question difficulty analysis,” *Array*, vol. 28, no. December, p. 100523, Dec. 2025, doi: 10.1016/j.array.2025.100523.

[13] H. Nakagawa, Y. Iwasawa, and Y. Matsuo, “Graph-based knowledge tracing: Modeling student proficiency using graph neural networks,” *Web Intelligence*, vol. 19, no. 1–2, pp. 87–102, Dec. 2021, doi: 10.3233/WEB-210458.

[14] A. M. Olney, “Extraction of Concept Maps from Textbooks for Domain Modeling,” in Intelligent Tutoring Systems, vol. 6095, V. Aleven, J. Kay, and J. Mostow, Eds., in Lecture Notes in Computer Science, vol. 6095. , Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 390–392. doi: 10.1007/978-3-642-13437-1_80.

[15] A. Klašnja-Milićević, B. Vesin, M. Ivanović, and Z. Budimac, “E-Learning personalization based on hybrid recommendation strategy and learning style

identification," *Computers & Education*, vol. 56, no. 3, pp. 885–899, Apr. 2011, doi: 10.1016/j.compedu.2010.11.001.

[16] D. B. Lowe, A. J. Bucknell, and R. G. Webby, "Improving hypermedia development: a reference model-based process assessment method," in *Proceedings of the tenth ACM Conference on Hypertext and hypermedia : returning to our diverse roots: returning to our diverse roots, Darmstadt Germany: ACM*, vol. 1999, no. February, pp. 139–146, 1999, doi: 10.1145/294469.294507.

[17] R. Ferguson, "Learning analytics: Drivers, developments and challenges," *International Journal of Technology Enhanced Learning*, vol. 4, no. 5–6, pp. 304–317, 2012, doi: 10.1504/IJTEL.2012.051816.

[18] A. T. Corbett and J. R. Anderson, "Knowledge tracing: Modeling the acquisition of procedural knowledge," *User Modeling and User-Adapted Interaction*, vol. 4, no. 4, pp. 253–278, 1995, doi: 10.1007/BF01099821.

[19] S. Minn, Y. Yu, M. C. Desmarais, F. Zhu, and J.-J. Vie, "Deep Knowledge Tracing and Dynamic Student Classification for Knowledge Tracing," in *2018 IEEE International Conference on Data Mining (ICDM), Singapore: IEEE*, vol. 2018, no. December, pp. 1182–1187, 2018, doi: 10.1109/ICDM.2018.00156.

[20] H. Drachsler and W. Greller, "Privacy and analytics: it's a DELICATE issue a checklist for trusted learning analytics," in *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge - LAK '16, Edinburgh, United Kingdom: ACM Press*, vol. 2016, no. April, pp. 89–98, 2016, doi: 10.1145/2883851.2883893.

[21] K. Verbert et al., "Context-Aware Recommender Systems for Learning: A Survey and Future Challenges," *IEEE Trans. Learning Technol*., vol. 5, no. 4, pp. 318–335, Oct. 2012, doi: 10.1109/TLT.2012.11.

[22] J. Portisch, N. Heist, and H. Paulheim, "Knowledge graph embedding for data mining vs. knowledge graph embedding for link prediction – two sides of the same coin?," *SW*, vol. 13, no. 3, pp. 399–422, Apr. 2022, doi: 10.3233/SW-212892.

[23] J. Zhang, N. Lin, X. Zhang, W. Song, X. Yang, and Z. Peng, "Learning Concept Prerequisite Relations from Educational Data via Multi-Head Attention Variational Graph Auto-Encoders," in *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining, Virtual Event AZ USA: ACM*, vol. 2022, no. February, pp. 1377–1385, 2022, doi: 10.1145/3488560.3498434.

[24] Z. Wu et al., "A comprehensive survey on graph neural networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4–24, 2021, doi: 10.1109/TNNLS.2020.2978386.

[25] Z. Wu, L. Huang, Q. Huang, C. Huang, and Y. Tang, "SGKT: Session graph-based knowledge tracing for student performance prediction," *Expert Systems with Applications*, vol. 206, no. November, p. 117681, Nov. 2022, doi: 10.1016/j.eswa.2022.117681.

[26] Q. Huang and Y. Zeng, "Improving academic performance predictions with dual graph neural networks," *Complex Intell. Syst*., vol. 10, no. 3, pp. 3557–3575, Jun. 2024, doi: 10.1007/s40747-024-01344-z.

[27] S. Hakkal and A. A. Lahcen, "Leveraging graph neural network for learner performance prediction," *Expert Systems with Applications*, vol. 293, no. December, p. 128724, Dec. 2025, doi: 10.1016/j.eswa.2025.128724.