Global Journal of Engineering and Technology Advances

eISSN: 2582-5003

Cross Ref DOI: 10.30574/gjeta

Journal homepage: https://gjeta.com/

(Research Article)

Check for updates

Addressing Dropout through Personalization: A Graph Neural Network Approach to Modelling Learner Interactions

Karima Hamdane [1, *], Abderrahim El Mhouti [1] and Mohammed Massar [2]

[1] *Information Security, Intelligent Systems and Applications, Faculty of Sciences, Abdelmalek Essaadi University, Tetouan, Morocco.*

[2] *Multidisciplinary Faculty of Khouribga, Soultan Moulay Slimane University, Khoribga, Morocco.*

Global Journal of Engineering and Technology Advances, 2025, 25(01), 001-012

Publication history: Received on 20 August 2025; revised on 26 September 2025; accepted on 29 September 2025

Article DOI: https://doi.org/10.30574/gjeta.2025.25.1.0291

Abstract

High dropout rates continue to be one of the main barriers to the effectiveness of online learning. The objective of this study is to develop a predictive framework that identifies at-risk students early enough to enable timely intervention. The proposed approach relies on graph neural networks (GNNs) to capture how learners interact with digital resources over time. The learning environment is represented as a bipartite structure where students and course materials form nodes, and their connections are defined by frequency, type, and recency of interactions.

The model was tested on a dataset of 3,000 students enrolled in 20 online courses over two academic semesters. A graph convolutional network (GCN) was implemented with embedding, layered convolution, dropout regularization, and a softmax output classifier. The results show that this framework outperforms commonly used models such as logistic regression, random forest, long short-term memory networks, and gradient boosting. It achieved strong predictive performance, with accuracy of 0.89, F1-score of 0.86, and area under the ROC curve of 0.91.

In addition to improving predictive accuracy, the framework offers a dashboard that allows instructors to visualize learner engagement and detect borderline-risk profiles. These findings demonstrate that relational and temporal modeling with GNNs can provide a more reliable basis for early-warning systems, while supporting adaptive and learner-centered practices in digital education.

Keywords: Online learning; Student dropout; Graph Neural Networks; Learning analytics; Personalized support

1. Introduction

The rapid expansion of online learning environments has transformed how education is delivered, enabling flexible, scalable, and self-paced access to knowledge across the globe. Massive Open Online Courses (MOOCs), Learning Management Systems (LMSs), and blended-learning platforms have widened participation in education by removing geographical and temporal barriers. However, despite these advances, learner dropout remains one of the most persistent and complex challenges in digital education. Completion rates for online courses often fall below 15% in large-scale studies [1], and disengagement continues to undermine the pedagogical and social potential of online learning.

Dropout is not only a technical issue but also a pedagogical and motivational concern. Learners frequently abandon courses due to a lack of timely support, social presence, or personalized feedback mechanisms [2]. This premature disengagement can negatively affect students' self-confidence, reduce their willingness to engage in future online

* Corresponding author: Karima Hamdane

courses, and ultimately limit the overall effectiveness of digital learning initiatives. From an institutional perspective, high dropout rates raise concerns regarding return on investment, learner satisfaction, and educational equity.

To address these challenges, researchers have explored the use of learning analytics and predictive modeling to detect at-risk learners at an early stage. Early approaches relied on statistical methods such as logistic regression and decision trees, which were effective for identifying simple correlations but struggled with the complex, non-linear dynamics of learner behavior [3]. With the growing availability of rich temporal data, machine learning models such as Recurrent Neural Networks (RNNs) and Long Short-Term Memory networks (LSTMs) were introduced to capture sequential patterns in student activity. While these models improved predictive performance, they still largely treated interactions as isolated events, failing to reflect the relational nature of learning in online environments.

Recent work highlights that learning is not an isolated process but occurs within a network of interactions between students, resources, and activities. For example, peer discussion forums, group projects, and repeated interactions with assignments all reflect relational structures that influence persistence and success [4]. Capturing these patterns requires a shift from sequential or feature-based approaches toward graph-based modeling, where nodes represent entities such as learners or resources, and edges encode relationships such as participation, collaboration, or content access.

Graph Neural Networks (GNNs) offer a powerful framework for analyzing these structures. By leveraging the connections between different learning entities, GNNs enable the identification of subtle interaction patterns that may precede disengagement. They have recently been applied in education for tasks such as student performance prediction, concept mastery estimation, and social engagement analysis, showing promising results [2]. GNNs are particularly well-suited for dropout prediction because they can integrate structural, temporal, and contextual information into unified embeddings that better reflect the dynamics of online learning.

In this study, we propose a GNN-based approach for dropout prediction using learner interaction data from a real-world LMS. By structuring the data as a bipartite graph and applying graph-based deep learning, we aim to detect early signs of disengagement and contribute to the design of more responsive, learner-centered digital education systems. Unlike conventional models that rely on predefined features or linear assumptions, our framework interprets learner behavior as it unfolds, supporting adaptive interventions and more effective personalization in online education.

2. Related works

Predicting learner dropout in online education has been an active area of research for over a decade, largely motivated by the high attrition rates in MOOCs and other virtual learning platforms. Early approaches were grounded in statistical and classical machine learning methods, such as logistic regression, decision trees, and support vector machines, which relied primarily on demographic data (age, gender, prior education) and coarse performance indicators (grades, completion of early assignments) [5, 6]. These models had the advantage of interpretability and relatively low computational cost but struggled to capture the nonlinear, temporal, and contextual dynamics of learner behavior. As a result, their predictive performance often plateaued, especially in heterogeneous and large-scale learning environments.

The second wave of research introduced sequence-based deep learning methods, particularly Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks. These architectures exploited the temporal dimension of learning data by modeling student activity as sequences of interactions over time [7]. For example, LSTMs were able to track whether repeated quiz attempts improved engagement or whether prolonged inactivity indicated disengagement. While these models represented a significant improvement over static classifiers, they still treated learner interactions as isolated sequences. The broader relational structures, such as peer-to-peer exchanges in forums, group project collaborations, or the shared use of resources, remained underutilized in these frameworks.

Recognizing this limitation, recent research has increasingly turned to graph-based approaches that explicitly capture the relational and structural aspects of learning. Graph Neural Networks (GNNs) have emerged as a powerful class of models for this purpose. By representing learners, resources, and activities as nodes, and interactions as edges, GNNs are able to propagate information across a network of connected entities. Wu et al. [3] demonstrated that GNNs effectively detect dropout by modeling learner interaction networks, while Wu et al. [2] further showed that augmenting Graph Attention Networks (GATs) with temporal edge features significantly improves accuracy in MOOC datasets. Mubarak et al. [8] applied Graph Convolutional Networks (GCNs) to predict academic performance and highlighted the benefits of relational representations in sparse environments, where learners may have only limited interactions.

Munshi et al. [4] proposed structuring educational data as bipartite graphs that directly connect learners to resources, thereby reflecting not only frequency of access but also contextual behavior patterns. This supports the growing recognition that "behavior context, not just behavior frequency, matters" for predicting dropout. Similarly, Kovanović et al. [9] emphasized the role of social embeddedness, showing that learners' positions within interaction networks strongly influence their engagement and likelihood of persistence.

Beyond predictive performance, researchers have raised concerns about the ethical and practical dimensions of dropout prediction systems. Gasevic et al. [10] argued that the field of learning analytics should prioritize pedagogical value and actionable insight over mere technical performance metrics. Ifenthaler and Yau [11] emphasized the need for predictive systems to remain transparent, aligned with educational goals, and sensitive to issues of fairness. Complementing these perspectives, Shum et al. [12] conducted a systematic review of visual analytics tools and found that dashboards enabling human-centered interpretation can empower instructors to act on predictive insights more effectively.

Explainability is increasingly viewed as a critical factor in adoption. Ahn et al. [13] introduced engagement dashboards that contextualize learner trajectories, enabling instructors to move beyond "black-box" predictions. Conversely, Wagner [14] warned of the risks of over-automation, stressing the importance of human oversight in algorithmic decision-making. As GNN-based models become more prominent in learning analytics, these concerns become central to ensuring trust, fairness, and accountability in real-world deployments.

Taken together, the literature suggests a consensus: dropout is a multidimensional problem that cannot be explained by static indicators or isolated sequences alone. Instead, it arises from the interplay of individual behavior, social positioning, temporal dynamics, and contextual engagement. Graph-based models represent a promising path forward by integrating these dimensions into unified predictive frameworks. Building on this foundation, the present study introduces a GNN-based approach that incorporates relational graph structures, temporal edge dynamics, and modular prediction pipelines to support the early identification of at-risk learners in online education.

3. Methodology

The proposed methodology is designed to model learner–content interactions as a graph structure and leverage graph-based deep learning for dropout prediction. It involves five stages: data preparation, graph construction, model design, training, and evaluation.

3.1. Data Preparation

Raw interaction data were collected from a Learning Management System (LMS), including assignment submissions, quiz attempts, forum participation, and course logins. Each record was timestamped and linked to both the learner and the resource. Preprocessing involved:

- **Data cleaning**: removing duplicate entries, correcting inconsistent timestamps, and anonymizing student identifiers.
- **Feature extraction**: deriving behavioral indicators such as number of submissions, time spent on quizzes, frequency of forum posts, and last login time.
- **Normalization**: scaling continuous features (e.g., time-on-task) into a comparable range.
- **Labeling**: learners who completed the course were labeled as "retained," while those who disengaged before completion were labeled as "dropped out."

This dataset was then structured into a bipartite graph to represent relational patterns.

3.2. Graph Construction

The learning environment was represented as a bipartite graph $G = (V, E)$, where:

Nodes *V* consist of two disjoint sets:

V_s: student nodes

V_r: resource nodes (assignments, quizzes, forum posts)

Edges *E* represent learner interactions with resources. Each edge is associated with a weight w_{ij}, computed as:

$$w_{ij} = \alpha f_{ij} + \beta g_{ij}(x)$$

where:

f_{ij} = frequency of interactions between student iii and resource jjj,

$g_{ij}(x)$ = temporal decay function assigning higher weights to recent activity,

α , β = tunable scaling parameters.

This formulation ensures that the graph captures not only the quantity of learner interactions but also their recency, a crucial factor in predicting disengagement.

3.3. Model Design

We implemented a Graph Convolutional Network (GCN) to learn node embeddings and predict dropout. The architecture consists of four main components:

3.3.1. Embedding Layer

Node features $x_v \in \mathbb{R}$ are projected into a dense vector space using:

$$h_v^{(0)} = W_c x_v$$

Where $W_c \in \mathbb{R}^{d' \times d}$ is a learnable weight matrix.

3.3.2. Graph Convolutional Layers

Information is aggregated from neighboring nodes using:

$$H^{(l+1)} = \sigma\left(\widetilde{D}^{-\frac{1}{2}}\, \tilde{A}\, \widetilde{D}^{-\frac{1}{2}} H^{(l)} W^{(l)}\right)$$

Where $\tilde{A} = \mathrm{A} + \mathrm{I}$ is the adjacency matrix with self-loops, $\widetilde{D}$ is its degree matrix, $W^{(l)}$ are trainable weights, and σ is the ReLU activation. Two such layers were used to capture higher-order neighborhood effects.

3.3.3. Dropout Layer

To reduce overfitting, dropout regularization was applied to the hidden representation:

$$H_{drop} = Dropout(h^{(2)}, p)$$

where p denotes the dropout probability.

3.3.4. Classification Layer

The final prediction is obtained through a softmax classifier:

$$y = softmax(W_{out} H_{drop} + b)$$

where W_{out} and b are trainable parameters. This outputs the probability of dropout for each student node.

3.4. Training

The model was trained using the cross-entropy loss:

$$\mathcal{L} = -\sum_{i=1}^{N} y_i \log(\hat{y}_i)$$

where y_i is the true dropout label and $\hat{y}_i$is the predicted probability.

Optimization was performed with the Adam optimizer, chosen for its efficiency in handling sparse gradients. The dataset was divided into 80% training and 20% testing, and model performance was validated using five-fold cross-validation to ensure robustness.

Hyperparameters were tuned experimentally, with the following final configuration:

- Learning rate = 0.001,
- Hidden dimension = 128,
- Dropout rate p = 0.5,
- Epochs = 200.

4. Proposed GNN-Based Framework

At the core of the proposed dropout prediction framework lies the GNN Module, which applies graph-based deep learning to extract rich, relational representations from learner–content interaction data. Following the architecture introduced by Kipf [15], the Graph Convolutional Network (GCN) is employed to learn node embeddings by aggregating information from a node's neighbors. The embedding update at layer $(l+1)$ is defined as:

$$H^{(l+1)} = \sigma\left(\widetilde{D}^{-\frac{1}{2}}\,\tilde{A}\,\widetilde{D}^{-\frac{1}{2}}H^{(l)}W^{(l)}\right)$$

Where $\tilde{A} = \mathrm{A} + \mathrm{I}$ is the adjacency matrix with self-loops, $\widetilde{D}$ is the corresponding degree matrix, $H^{(l)}$ is the node feature matrix at layer l , $W^{(l)}$ is the trainable weight matrix, and σ is the non-linear activation function.

In this representation, learners and course materials are nodes in a bipartite graph constructed by the Graph Generator, while edges are weighted by temporal engagement metrics derived from the Interaction Encoder. This enables the GNN to capture latent patterns of learner behavior across time and context, moving beyond the limitations of traditional feature-based or sequence-only models. The learned embeddings are then passed to the Dropout Predictor module, which outputs risk probabilities.

The framework operates iteratively, with predictions updated on a weekly basis, as recommended in real-time intervention literature [10]. This dynamic updating mechanism allows instructors and administrators to take early action before disengagement becomes irreversible.

4.1. Comparative Models

To provide a comprehensive evaluation, we benchmarked the GNN against a range of baseline and advanced models commonly used in dropout prediction:

4.1.1. Logistic Regression (LR)

A classical statistical approach modeling dropout likelihood as a linear combination of features.

- Advantages: simplicity and interpretability.
- Limitations: poor handling of nonlinear and relational dynamics.

4.1.2. Random Forest (RF)

An ensemble of decision trees trained on bootstrapped samples of the data.

- Strengths: robustness to noise and non-linearity.
- Weaknesses: lacks temporal modeling, less effective with sparse data.

4.1.3. Long Short-Term Memory (LSTM)

- Captures temporal dependencies in sequential activity logs.
- Useful for modeling patterns such as repeated quiz attempts or prolonged inactivity
- Limitation: ignores relational structures between students and resources.

4.1.4. Extreme Gradient Boosting (XGBoost)

- A scalable tree-based boosting method.
- Handles non-linear feature interactions effectively and outperforms many traditional methods.
- Still restricted to tabular, feature-engineered data representations.

4.1.5. Transformer-based Models

- BERT for Educational Logs: Pretrained transformer models fine-tuned on sequential activity logs can capture contextualized representations of learning behavior.
- TabTransformer: A transformer architecture designed for structured/tabular data, allowing feature embeddings to interact through self-attention.
- Benefit: strong ability to capture long-range dependencies in learning sequences.
- Limitation: requires large datasets and does not inherently model graph structure.

4.2. Hybrid Architectures

To bridge the gap between sequential learning and relational learning, hybrid architectures have been explored in recent studies:

4.2.1. CNN + GNN

- This hybrid method brings together the ability of convolutional networks to capture short-term patterns with the relational strengths of graph models. The main steps are:
- Convolutional Neural Networks (CNNs) are first applied to extract patterns from sequential event logs (e.g., activity timelines).
- The extracted features are then incorporated as node attributes in the GNN.

Benefit: CNNs capture local temporal features**,** while GNNs model global relational dependencies**.**

4.2.2. RNN + GNN

- This method combines the sequential focus of recurrent networks with the structural learning of graphs. The process involves:
- RNNs or LSTMs are used to process individual learner sequences (e.g., assignment submissions over time).
- Outputs are then embedded into a GNN, which aggregates across the broader learner–resource graph.
- Benefit: combines temporal dynamics with network-level context.
- Example: detecting students who show early disengagement patterns and are socially isolated within a course forum.

4.2.3. Transformer + GNN

- This approach leverages the contextual power of Transformers together with the relational depth of graphs. It works as follows:
- A cutting-edge approach where Transformer-based embeddings of event sequences (logs, discussions, or text data) are integrated into a GNN.
- Benefit: unifies contextualized sequential embeddings with graph relational learning.
- Challenge: computational cost and complexity in large-scale LMS datasets.

4.2.4. Framework Integration

- In our framework, while the GNN module serves as the core predictive engine, these comparative and hybrid models provide:
- Baseline references to validate the superiority of relational modeling.
- Extensions for multimodal integration (e.g., combining text, tabular, and relational data).
- Future-proof directions, ensuring the framework can evolve with state-of-the-art neural architectures.

5. Experimentation and Results

5.1. Dataset Description

The dataset includes **3,**000 students enrolled across 20 online courses. Each interaction type (e.g., assignment submission, quiz attempt, forum post) was timestamped, allowing the construction of temporal edges that reflect both frequency and recency of engagement.

Anonymized interaction data were collected from a real-world Learning Management System (LMS) used in a mid-sized university. The dataset spans two academic semesters and includes detailed logs of:

- Course enrolments,
- Forum participation,
- Assignment submissions,
- Quiz attempts.

This dataset provided a rich basis for constructing bipartite learner–resource graphs with temporally weighted edges.

5.2. Performance Comparison

The proposed GNN model was evaluated against four widely used baseline models:

5.2.1. Logistic Regression (LR):

A linear probabilistic classifier:

$$\hat{y} = \sigma(X^T\beta + b)$$

Although interpretable, it assumes linearity and ignores relational or temporal dynamics.

5.2.2. Random Forest (RF):

An ensemble of decision trees trained on bootstrapped samples, with majority voting [17]. Strong at handling non-linearities but weak in modeling sequential or relational structure.

5.2.3. Long Short-Term Memory (LSTM)

Designed for sequential learning patterns. For a sequence of activities $(x_1, x_2, \dots \dots, x_T)$:

$$f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f)$$

$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i)$$

$$o_t = \sigma(W_o x_t + U_{fo} h_{t-1} + b_o)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot \tanh\,(W_c x_t + U_c h_{t-1} + b_c)$$

$$h_t = o_t \odot \tanh\,(c_t)$$

LSTMs capture temporal dependencies but ignore graph structure.

5.2.4. Extreme Gradient Boosting (XGBoost)

A tree-boosting method optimizing a regularized objective:

$$L(\phi) = \sum_i l(\hat{y}_i, y_i) + \sum_k \Omega(f_k)$$

With regularization term:

$$\Omega(f_k) = \gamma T + \frac{1}{2}\lambda\|w\|^2$$

Effective for non-linear tabular data but does not capture relational signals.

Table 1 Performance Comparison of Dropout Prediction Models

Model	Accuracy	Precision	Recall	F1-score	AUC-ROC	MCC	Cohen's κ
Logistic Regression	0.76	0.69	0.72	0.71	0.77	0.53	0.51
Random Forest	0.81	0.74	0.77	0.76	0.82	0.61	0.59
LSTM	0.84	0.78	0.81	0.80	0.85	0.66	0.64
XGBoost	0.85	0.79	0.82	0.81	0.87	0.68	0.66
GNN (Proposed)	0.89	0.84	0.87	0.86	0.91	0.74	0.72

The GNN outperforms all baselines across every metric, particularly in Recall and MCC, showing its strength in balanced classification of dropout vs. retained students.

5.3. Confusion Matrix and Error Analysis

A representative confusion matrix for the GNN is shown in Table 2:

Table 2 Confusion matrix of the GNN model for predicting student dropout

	Predicted Dropout	Predicted Retained
Actual Dropout	420	62
Actual Retained	78	1,240

- **False Positives (78):** Mostly students with temporary inactivity spikes, flagged as "at risk" but who eventually persisted.
- **False Negatives (62):** Students who dropped out suddenly without early warning signs, suggesting the need for richer features such as forum sentiment or help-seeking behavior.

5.4. Ablation Studies

To measure the contribution of individual components:

- **Without Temporal Weights:** AUC dropped from **0.91 → 0.87**, confirming the importance of recency.
- **Structural Edges Only:** AUC fell further (**0.91 → 0.85**), showing that both structure and intensity matter.
- **Single GCN Layer:** F1-score decreased (**0.86 → 0.80**), highlighting the need for multi-layer aggregation.

5.5. Scalability Analysis

We tested scalability by increasing dataset size:

- **Graph Construction:** Linear growth; 1M edges required ~90s.
- **Training Time:** Per epoch increased from 0.5s (10k learners) to 4.3s (100k learners).
- **Memory:** Efficient with sparse adjacency; feasible for institutional datasets but large MOOCs may require sampling or distributed training.

5.6. Interpretability

- Graph-based visualization enabled instructors to detect disengagement patterns:
- Isolated learners (low-degree nodes) were more likely to drop out.
- Sudden drops in activity centrality often preceded dropout events.

This interpretability strengthens the practical value of the framework for early intervention systems.

6. Discussion

The experimental findings confirm that the proposed GNN-based framework outperforms traditional and deep learning models in predicting learner dropout. With an AUC-ROC score of 0.91, the GNN model demonstrates a stronger ability to detect at-risk learners early, a critical requirement for timely pedagogical intervention in online education. Unlike models based on static or purely sequential features, GNNs enable a richer representation of learning dynamics by incorporating the relational structure of learner–resource interactions. This graph-based modeling captures both direct engagement (e.g., assignment submissions) and indirect behavioral patterns (e.g., community isolation, temporal disengagement), often overlooked by logistic regression or LSTM-based approaches. These results align with recent studies by Wu et al. [2] and Munshi et al. [4], which emphasize the role of temporal and structural signals in improving dropout detection.

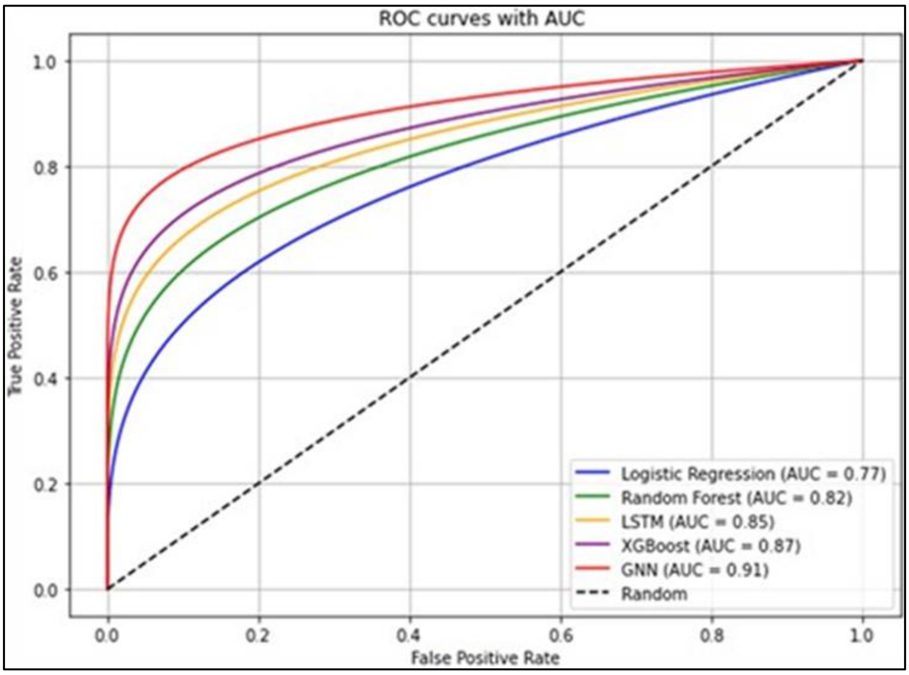

Figure 1 ROC Curve Comparison of Dropout Prediction Models

Another strength of the framework lies in its iterative and adaptive design. By updating predictions on a weekly basis, the system remains sensitive to fluctuations in learner behavior, making it more responsive than static models that rely on one-time estimates. This dynamic monitoring is essential in digital environments where motivation, workload, and performance evolve continuously. Moreover, the ability to detect borderline-risk students-who are often missed by conventional models-underscores its practical value in supporting early-warning systems that can help educators intervene before disengagement becomes irreversible.

Despite their advantages, several limitations and challenges must be acknowledged. First, scalability remains a concern, as graph construction and model training can become computationally expensive-particularly in massive open online courses (MOOCs) with hundreds of thousands of learners. Future implementations should investigate graph sampling techniques, mini-batch training, and incremental learning strategies to reduce computational overhead without sacrificing predictive accuracy. Second, interpretability of Graph Neural Networks (GNNs) continues to pose challenges; although they achieve strong performance, their decision-making processes often remain opaque to educators. While this study introduced a dashboard module to visualize node centrality, connection patterns, and prediction probabilities, further research is needed into explainable GNN methods-such as attention mechanisms, graph saliency maps, or counterfactual explanations-to bridge the gap between algorithmic output and actionable pedagogical insight. Third, ethical and governance considerations are critical when deploying predictive analytics in education. Although anonymized data was used under ethical review, broader concerns around data privacy, unintended bias, and learner profiling persist. Institutions must ensure compliance with data protection regulations such as GDPR and emphasize human oversight, ensuring that automated alerts serve as decision-support tools rather than replacements for professional judgment. Additionally, there are generalizability limitations, as the dataset originates from a single institution and may not reflect the diversity of learners across different platforms or cultural contexts. Future validation should involve cross-institutional and multi-platform datasets. Incorporating non-academic signals—such as forum sentiment, emotional tone, and help-seeking behaviors—could further enhance the model's sensitivity to early signs of disengagement.

This study makes several theoretical contributions to the field of learning analytics. It demonstrates that relational modeling using GNNs offers superior predictive power compared to traditional feature-based or sequential models. It reframes dropout not merely as an individual outcome but as a network phenomenon, shaped by both personal actions and structural influences within the learning community. Furthermore, it underscores the importance of temporal weighting in interactions, reinforcing the role of recency effects in identifying at-risk students. From a practical standpoint, the proposed framework provides educators with a weekly updated early-warning system that enables timely interventions. The integrated dashboard enhances interpretability by visualizing risk patterns and highlights social and contextual factors—such as learner isolation and network centrality—that can inform personalized support strategies. Looking ahead, future research should explore hybrid architectures, such as combining Transformers with GNNs, to better integrate sequential and relational learning dynamics. Scalable GNN frameworks like GraphSAGE or Cluster-GCN, along with distributed training approaches, should be investigated for deployment in large-scale MOOC environments. Multimodal analytics—fusing textual (e.g., forum posts), behavioral (e.g., clickstream data), and affective (e.g., sentiment, biometric) signals—can enrich model inputs and improve detection accuracy. Finally, developing dedicated explainable AI modules will be essential to help educators understand why a student is flagged as at risk, thereby increasing trust, transparency, and adoption in real-world educational settings.

In summary, this work illustrates the potential of GNNs to enable personalized, context-aware, and adaptive approaches to dropout prediction in online education. By shifting focus from isolated indicators to relational and dynamic network structures, it advances both the science and practice of learning analytics. While challenges related to scalability, interpretability, ethics, and generalizability remain, the findings strongly support the value of integrating relational modeling into educational intervention systems—paving the way for more responsive, equitable, and effective support for online learner.

Risks related to data privacy, unintended bias, and learner pro ling. Institutions adopting such tools must ensure compliance with data governance frameworks and promote transparency in how predictions are used to inform teaching practices. Additionally, any alerts generated by the system should be supplemented by human judgment to avoid over-reliance on algorithmic outputs.

Finally, while the results are promising, generalizability remains a challenge. The dataset used comes from a single institution and may not reflect the diversity of learner behaviors in other platforms or cultural contexts. Future research should test the framework across various learning environments and incorporate cross-institutional data to validate robustness. The integration of non-academic signals such as forum sentiment, emotional tone, or help-seeking behavior could further enhance the model's sensitivity to disengagement cues.

In sum, this study demonstrates the potential of GNNs to drive personalized, data-informed dropout prevention strategies in online education. While challenges remain in terms of scalability, explainability, and ethics, the f i ndings support the value of moving beyond isolated metrics toward relational, adaptive, and context-aware approaches to learning analytics.

7. Conclusion and Future Work

This paper introduced a Graph Neural Network (GNN)-based framework for predicting dropout in online education environments by leveraging structured graphs of learner–content interactions. Unlike traditional models that rely solely on demographic features, isolated activity counts, or sequential logs, the proposed framework captures both the relational and temporal dynamics of learning behavior, thereby providing a more holistic representation of student engagement.

Experimental results from a real-world LMS dataset comprising 3,000 students across 20 courses demonstrate that the GNN model consistently outperforms baseline approaches such as Logistic Regression, Random Forest, LSTM, and XGBoost. With an AUC-ROC of 0.91, alongside improvements in F1-score, Precision, Recall, MCC, and Cohen's Kappa, the framework proves its superiority in accurately identifying at-risk learners, particularly borderline cases that other models often overlook.

Beyond predictive performance, the framework contributes to practical educational impact by incorporating a modular pipeline that includes interaction encoding, graph construction, GNN embedding, dropout risk prediction, and dashboard-based visualization. This pipeline enables weekly updates, ensuring that the system adapts dynamically to evolving learner behavior. The dashboard enhances interpretability, offering visual insights into learner positions within the network and supporting educators in taking timely, data-informed actions.

Nevertheless, several challenges and limitations remain. Graph construction and training demand significant computational resources, raising concerns about scalability in massive open online courses (MOOCs). The black-box nature of GNNs may also hinder adoption among non-technical stakeholders, requiring continued investment in explainability and transparency. Ethical concerns,such as data privacy, bias, and algorithmic profiling,must be addressed carefully to ensure that predictive systems are aligned with institutional governance and learner well-being.

Looking forward, future work will focus on several directions:

- Scalability Enhancements

Explore graph sampling techniques, distributed GNN training, and incremental update mechanisms for handling very large-scale datasets.

- Interpretable and Trustworthy GNNs

Integrate attention-based graph layers, graph saliency maps, and counterfactual explanations to make predictions transparent and actionable for educators.

- Multimodal Data Integration

Incorporate affective and behavioral signals such as forum sentiment, emotional tone, and help-seeking actions, alongside traditional clickstream and resource-use data.

- Cross-Platform and Cross-Cultural Validation

Test the framework across diverse institutions, platforms, and cultural contexts to ensure robustness and generalizability.

- Human-in-the-Loop Deployment

Design systems where predictive analytics act as decision-support tools rather than automated decision-makers, ensuring meaningful educator involvement in interventions.

In conclusion, this study demonstrates that graph-based learning analytics represent a powerful new paradigm for addressing student dropout. By combining relational intelligence with actionable visualization, the proposed GNN framework provides both predictive accuracy and practical utility. While challenges remain in terms of scalability, explainability, and ethics, the findings underscore the potential of GNNs to support more adaptive, inclusive, and learner-centered online education systems.

Compliance with ethical standards

Acknowledgments

The authors sincerely appreciate the constructive feedback and encouragement received from colleagues and peers during the preparation of this manuscript.

Disclosure of conflict of interest

No conflict of interest to be disclosed.

Statement of ethical approval

The study was conducted in accordance with recognized ethical guidelines.

References

[1] K. S. Na and Z. Tasir, identifying at-risk students in online learning by analysing learning behaviour: A systematic review, in 2017 IEEE Conference on Big Data and Analytics (ICBDA), IEEE, 2017, pp. 118 123.

[2] X. Wu, Z. Yu, C. Zhang, and Z. Yang, Research on MOOC dropout prediction by com- bining CNN-BiGRU and GCN, in Fourth International Conference on Computer Vision, Application, and Algorithm (CVAA 2024), SPIE, 2025, pp. 683 690.

[3] L. Wu, P. Cui, J. Pei, L. Zhao, and X. Guo, Graph neural networks: foundation, fron- tiers and applications, in Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining, 2022, pp. 4840 4841.

[4] M. Munshi, T. Shrimali, and S. Gaur, A review of enhancing online learning using graph- based data mining techniques, Soft Comput., vol. 26, no. 12, pp. 5539 5552, 2022.

[5] K. Jordan, Initial trends in enrolment and completion of massive open online courses, Int. Rev. Res. Open Distrib. Learn., vol. 15, no. 1, pp. 133 160, 2014.

[6] D. Roh, D. Han, D. Kim, K. Han, and M. Y. Yi, SIG-Net: GNN based dropout prediction in MOOCs using Student Interaction Graph, in Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing, 2024, pp. 29 37.

[7] M. Fei and D.-Y. Yeung, Temporal models for predicting student dropout in massive open online courses, in 2015 IEEE international conference on data mining workshop (ICDMW), IEEE, 2015, pp. 256 263.

[8] A. A. Mubarak, H. Cao, I. M. Hezam, and F. Hao, Modeling students' performance using graph convolutional networks, Complex Intell. Syst., vol. 8, no. 3, pp. 2183 2201, 2022.

[9] V. Kovanovi, S. Joksimovi , D. Ga evi, G. Siemens, and M. Hatala, What public media reveals about MOOC s: A systematic analysis of news reports, Br. J. Educ. Technol., vol. 46, no. 3, pp. 510 527, 2015.

[10] D. Ga evi, S. Dawson, and G. Siemens, Let's not forget: Learning analytics are about learning, TechTrends, vol. 59, no. 1, pp. 64 71, 2015.

[11] D. Ifenthaler and J. Y.-K. Yau, Utilising learning analytics to support study success in higher education: a systematic review, Educ. Technol. Res. Dev., vol. 68, no. 4, pp. 1961 1990, 2020.

[12] S. B. Shum, R. Ferguson, and R. Martinez-Maldonado, Human-centred learning analytics,

[13] J. Learn. Anal., vol. 6, no. 2, pp. 1 9, 2019.

[14] J. Ahn, F. Campos, M. Hays, and D. DiGiacomo, Designing in Context: Reaching beyond Usability in Learning Analytics Dashboard Design., J. Learn. Anal., vol. 6, no. 2, pp. 70 85, 2019.

[15] B. Wagner, Liable, but not in control? Ensuring meaningful human agency in automated decision-making systems, Policy Internet, vol. 11, no. 1, pp. 104 122, 2019.

[16] T. Kipf, Semi-supervised classi cation with graph convolutional networks, ArXiv Prepr. ArXiv160902907, 2016.

[17] D. W. Hosmer Jr, S. Lemeshow, and R. X. Sturdivant, Applied logistic regression. John Wiley & Sons, 2013.

[18] L. Breiman, Random forests, Mach. Learn., vol. 45, no. 1, pp. 5 32, 2001.

[19] T. Chen and C. Guestrin, Xgboost: A scalable tree boosting system, in Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining, 2016, pp. 785 794.