# Enhancing the Performance of Automated Grade Prediction in MOOC using Graph Representation Learning

Soheila Farokhi<sup>†</sup> Utah State University soheila.farokhi@usu.edu Aswani Yaramal<sup>†</sup> Utah State University aswani.yaramala@usu.edu

sity Nanning Normal University 1.edu jiangtao@nnnu.edu.cn

Jiangtao Huang

Muhammad Fawad Akbar Khan<sup>†</sup> Utah State University khan@usu.edu

Xiaojun Qi Utah State University xiaojun.qi@usu.edu Hamid Karimi<sup>†</sup> Utah State University hamid.karimi@usu.edu

Abstract-In recent years, Massive Open Online Courses (MOOCs) have gained significant traction as a rapidly growing phenomenon in online learning. Unlike traditional classrooms, MOOCs offer a unique opportunity to cater to a diverse audience from different backgrounds and geographical locations. Renowned universities and MOOC-specific providers, such as Coursera, offer MOOC courses on various subjects. Automated assessment tasks like grade and early dropout predictions are necessary due to the high enrollment and limited direct interaction between teachers and learners. However, current automated assessment approaches overlook the structural links between different entities involved in the downstream tasks, such as the students and courses. Our hypothesis suggests that these structural relationships, manifested through an interaction graph, contain valuable information that can enhance the performance of the task at hand. To validate this, we construct a unique knowledge graph for a large MOOC dataset, which will be publicly available to the research community. Furthermore, we utilize graph embedding techniques to extract latent structural information encoded in the interactions between entities in the dataset. These techniques do not require ground truth labels and can be utilized for various tasks. Finally, by combining entityspecific features, behavioral features, and extracted structural features, we enhance the performance of predictive machine learning models in student assignment grade prediction. Our experiments demonstrate that structural features can significantly improve the predictive performance of downstream assessment tasks. The code and data are available in https://github.com/ DSAatUSU/MOOPer\_grade\_prediction

Index Terms—MOOC, Grade Prediction, Graph Representation Learning, Machine Learning

# I. INTRODUCTION

Massive Open Online Courses (MOOCs) are open courses with open enrollment that aim to make educational material accessible to a large audience online and often for free regardless of location or background [1]. Over the last ten years, MOOCs have emerged as a new online learning trend. In addition to dedicated providers (e.g., Coursera and Udemy), many

<sup>†</sup>These authors have an affiliation with Data Science and Applications Lab at Utah State University.

renowned universities have launched numerous courses that have drawn an oversized enrollment. MOOCs play a crucial role in modern education by democratizing access to quality learning [2], promoting lifelong learning [3], and offering flexibility [4], cost-effectiveness [5], and customization [6]. In addition, they facilitate global collaboration, rapid dissemination of knowledge [7], and targeted skill development [8] while providing a platform for educators to experiment with innovative teaching methods. By expanding the reach of educational institutions beyond traditional borders, MOOCs foster a global community of learners and educators, ultimately transforming the education landscape and making learning more accessible, flexible, and personalized. Automated grade prediction in MOOCs is a vital task with far-reaching benefits for students, instructors, and course providers. Some of these benefits include:

- Personalized learning: Automated grade prediction can help identify a student's strengths and weaknesses, enabling the creation of personalized learning paths to address specific knowledge gaps [9, 10].
- □ **Early intervention:** By predicting a student's performance, instructors and course providers can intervene early and offer additional support or resources to help struggling students succeed [11, 12].
- Improved engagement: With insights into student performance, course providers can make data-driven decisions to improve course design, engagement, and retention [13].
- Enhanced feedback: Automated prediction systems can provide immediate feedback to students, which can help them understand their progress and make necessary adjustments in their learning strategies [13].
- ☐ Motivation: Providing students with their predicted grades can encourage them to take control of their learning and motivate them to improve their performance [14, 15].
- Performance tracking: Automated grade prediction enables tracking of student performance over time [16],

helping instructors identify trends and patterns that can inform future course design and delivery.

- Dropout prevention: Identifying students at risk of dropping out or under-performing early on, allows for targeted interventions to help retain students and improve overall course completion rates [17, 18].
- □ **Informed decision-making:** Grade prediction can help course providers evaluate the effectiveness of their courses, identify areas for improvement, and make evidence-based decisions about future course offerings [19].
- Scalability: Automated grade prediction can be especially beneficial in MOOCs, which often have large numbers of students [20] because it can provide a more consistent and scalable way to predict and monitor student performance across diverse groups.



Fig. 1: Visualizing the traditional approach used in prior prediction models compared to our graph representation.

Over the past decade, numerous automated grade prediction methods have emerged in the MOOC landscape, offering various benefits, as mentioned above. However, existing methods still face two critical shortcomings. Firstly, current methods have a large granularity of grade prediction tasks, often focusing on significant assignments or entire courses [21]. In contrast, MOOC systems frequently include shorter exercises that target specific skills, such as Python's 'for loop' and 'continue statement.' Predicting future grades for these smaller exercises would be highly beneficial. In this paper, we refer to such exercises as 'challenges'-See Section III. Secondly, most existing methods overlook the complex yet rich structural relationships between pertinent entities, such as students and courses, which are essential to the prediction task. To illustrate, consider the visualization in Figure 1, which demonstrates our challenge grade prediction study. Figure 1(a) shows a scenario where the data of each student-challenge pair is treated separately, whereas Figure 1(b) presents our approach, which conceptualizes the relationship between relevant entities as an interaction graph. This graph encodes the inherent links between students and challenges. We hypothesize that the interaction graph between entities offers rich structural information, which can be extracted and leveraged to enhance the predictive performance of the grade prediction model.

In this study, we address the above-mentioned shortcomings by focusing on the challenge prediction task, which offers higher granularity than course outcome prediction. First, we introduce a new dataset from a large MOOC provider in China, containing various entities (e.g., challenge, course, chapter) for thousands of students. We then construct an interaction graph between students and challenges as a bipartite graph, as seen in Figure 1(b), and extract salient and dense entity-level vector representations using advanced graph representation learning techniques. Specifically, we employ two powerful node embedding learning methods, node2vec [22] and DeepWalk [23]. One of the main benefits of these graph representation learning methods is their unsupervised training, using only the underlying graph structure without ground truth labels, such as challenge grades. To validate the effectiveness of knowledge graph representations, we develop machine learning prediction models and conduct extensive experiments with our dataset. Our results show that incorporating structural graph information can enhance the predictive performance of challenge grade prediction. Worth noting that existing methods often rely on different and often complicated data and features, such as student forums [15] or watched videos [11], which may only sometimes be readily available.

- We introduce a new MOOC dataset that includes various entities and their interactions for several thousand students, which will be publicly available to the research community.
- Contrary to current grade prediction methods in MOOCs, we focus on short, small, and particular exercises, referred to as challenges in our dataset.
- ➤ We build a bipartite graph between students and challenges, then extract unsupervised dense entity representations using advanced neural networks.
- We conduct extensive experiments, demonstrating the usability of the interaction graph and its resulting representations for challenge grade prediction.

**Remark.** In this paper, the term "performance" is used to describe two different concepts. First, it refers to the academic achievement of students, specifically their grades. Second, it pertains to the performance of machine learning algorithms, specifically their accuracy. The distinction between these two uses of "performance" is made clear from the context in which it is used.

The rest of this paper is organized as follows. First, in Section II, we briefly review related work. Next, in Section III, we introduce the dataset, followed by our grade prediction problem statement in Section IV. Section V describes the methodology, and Section VI includes experiments and discussions. Finally, we conclude the paper in Section VII and enumerate several possible future directions.

# II. RELATED WORK

Educational data mining is an emerging field that leverages computational approaches to analyze large-scale educational data [24, 25, 26, 27, 28, 29, 30]. In particular, computational techniques for performance prediction in MOOCs were developed over the past ten years as MOOCs gained popularity. In [16], the authors developed a personalized linear multiple regression model to predict a student's future performance for specific assessment activities within a MOOC. They extracted six types of features (session features, quiz-related features, video-related features, homework-related features, time-related features, and interval-based features) from MOOC server logs to identify learning behavior and study habits for different students. Ramesh et al. [31] developed two distinct probabilistic soft logic models to forecast student achievement. One model used behavioral, linguistic, and structural features, while the other treated learner engagement as a latent variable and linked observed features to one or more types of engagement. In [15], the authors first classified learners according to motivation into three groups: certification earning, video watching, and course sampling. They then used an SVM-based model to predict grades by classifying certification learners into two classes (may or may not obtain a certification). Brinton and Chiang [32] applied standard algorithms to predict whether a MOOC student would be Correct on First Attempt (CFA) in answering a question. They found that parsing click-stream data into summary quantities was useful for classifying CFA. They computed nine video-watching summary quantities, such as fraction spent, fraction completed, fraction played, number of pauses, fraction paused, average playback rate, standard deviation of playback rate, number of rewinds, and number of fast forwards. Robust CFA prediction provided insights into student learning and helped instructors address students collectively. A continuous assessment prediction problem could be viewed as knowledge tracing, which predicted student's future performance based on their past activity. In [11], a time series neural network was trained using lecture video watching clickstreams for predicting student grades in MOOCs. Piech et al. [33] applied RNN and LSTM networks to predict students' responses to exercises based on their past activity in MOOCs. Furthermore, [9] used LSTM, RNN, LR, and SVM to make individual grade predictions, with LSTM achieving the best performance with an average AUC score of 0.748. Some researchers proposed novel deep learning models to improve MOOC learner performance prediction. In [9], the authors developed a unified model that incorporated student demographics, forum activities, and learning behavior to predict students' assignment grades. They cast the prediction task as a binary classification problem, predicting whether the student's grade was ranked in the top 30% of all students. Kim et al. [34] used deep learning methods to predict real-time student performance. They proposed a new deep learning-based method (GritNet), built upon the bidirectional long short-term memory (BLSTM) to predict student performance, such as graduation prediction. The algorithm significantly improved prediction quality within the first few weeks of the student experience. In [25], the authors introduced Deep Online Performance Evaluation (DOPE) to predict a student's performance in a specific course. Their method first represented the online learning system as a knowledge graph and learned student and course embeddings from historical data using a relational graph neural network. Additionally, DOPE used an LSTM to harness student behavior data into a condensed encoding since the data naturally took on a sequential shape. In [35], the authors introduced a model using Graph Neural Networks called  $R^2GCN$  (GNNs). They first created a network of questions and interactions among students, and then they utilized the proposed model to predict students' performance. Specifically, they formulated the student performance prediction in interactive online question pools as a node classification problem on a large heterogeneous network consisting of questions, students, and the interactions between them, better capturing the underlying relationship between questions and students.

All the previous studies on grade prediction in MOOCs mainly focus on course-level outcomes or major assignments, but little attention is paid to predicting grades for short, small, and specific exercises, which are called challenges in this study. Additionally, most existing methods ignore the complex structural relationships between pertinent entities, such as students and courses. To address these limitations, this study proposes a new approach that constructs an interaction graph between students and challenges and leverages advanced graph representation learning techniques to extract salient and dense entity-level vector representations. The effectiveness of this approach is validated through machine learning prediction models and extensive experiments on a new MOOC dataset introduced in this study.

## III. DATASET

MOOPer is an extensive online open practice dataset provided by a prominent Chinese university. The dataset contains students' online practice data from 2018 to 2019, along with supplementary information on users and practice projects organized in the form of knowledge maps. The MOOPer dataset is divided into two components: interactive data and knowledge maps. The interactive dataset comprises 2.53 million practice records, 21.6 million system feedback records, and 15,000 forum chat records. In the knowledge map section, there are 11 unique entity categories and ten specific relationship types.

### A. Knowledge map

The knowledge map provides information about entities and their relations. The structure of the knowledge map is illustrated in the left section of Figure 2.

1) Entities: The dataset includes entities such as discipline, sub-discipline, department, courses, chapter, challenge, exercise, school, student, teacher, and topic. More details about each entity will be provided in the dataset description upon its release.

2) *Relation:* This section explains the relationships between different entities. There is a hierarchical relationship among the entities. For instance, each course contains numerous exercises, and each exercise comprises multiple challenges. However, a single exercise can be linked to multiple courses.



Fig. 2: Entities and their relations in the MOOPer dataset.

# B. Interactive data

Student engagement with learning materials is categorized into three groups: user behavior, user feedback, and system feedback, described below.

1) User behavior: This category includes information on the challenges students completed and the grades they achieved for those tasks. It also contains data on the interaction's open and close time, the number of attempts to solve the challenge, whether students referred to the answers, and other additional information.

2) User feedback: This category displays the user's rating of the interactive practice activities and their choice of learning resources with different difficulty levels and problem types. Moreover, students can share their discussions in the forum. Chat content can be used to assess students' learning states and satisfaction, while the question-and-answer conversation can reveal their "blind spots" in knowledge acquisition. Additionally, the user's behavior in the forum serves as a critical indicator to speculate on their psychological condition and learning style [36].

3) System feedback: This category provides students with feedback on the system's results, such as syntax errors in their code, compilation results of the submitted code, and the discrepancy between the actual output and the intended output. Such data can be used to better understand users' learning abilities and knowledge mastery [37].

#### **IV. PROBLEM STATEMENT**

Suppose we have a subset of m challenges from our dataset denoted as  $C = \{c_1, c_2, \cdots, c_m\}$ . Also, let there be n students who have interacted with (or attempted) at least two of the *m* challenges in C, denoted as  $S = \{s_1, s_2, \cdots, s_n\}$ . For each challenge  $c_j$ , we assume the features can be represented as the vector  $\mathbf{f}_i \in \mathbb{R}^{d_c}$ , where  $d_c$  is the dimension size after encoding the challenge features. All challenge features are denoted as  $\mathcal{F} \in \mathbb{R}^{\bar{d}_c \times |\mathcal{C}|}$ . Similarly, for each student  $s_i$ , we assume the collected information can be represented as the vector  $\mathbf{d}_i \in \mathbb{R}^{d_s}$ , with  $d_s$  being the dimension size after encoding the student-related data. All student features are denoted as  $\mathcal{D} \in \mathbb{R}^{d_s \times |\mathcal{S}|}$ . In addition to the student-specific and challenge-specific data, the system is assumed to have collected some student-challenge data for each student  $s_i$  engaging with challenge  $c_j$ , which we represent as  $\mathbf{b}_{ij}$ . Let  $g_{ij}$  denote the grade of student  $s_i$  in

challenge  $c_j$ . Let  $\mathcal{B}$  and  $\mathcal{G}$  represent the entire studentchallenge data and grades, respectively. Moreover, suppose graph BG =  $(\mathcal{S}, \mathcal{C}, E)$  denotes the bipartite graph of studentchallenge interactions, where  $E \subset \mathcal{S} \times \mathcal{C}$  and  $e = (s_i, c_j) \in$ E is the edge between student  $s_i$  and challenge  $c_j$ . Finally, let  $\mathcal{M}(.)$  denote the graph representation machinery that takes graph BG and provides dense representations for both student and challenge pairs  $s_i$  and  $c_j$  in E.

Given the notations listed above, our goal is to learn a machine learning model f(.) that can predict students' challenge outcomes  $\mathcal{G}$  as follows:

$$f((\mathcal{C}, \mathcal{S}, \mathcal{F}, \mathcal{D}, \mathcal{B}, \mathcal{M}(\mathrm{BG})), \mathcal{G}) \to \hat{\mathcal{G}}$$

where  $\hat{\mathcal{G}}$  should be as close as possible to  $\mathcal{G}$ .

# V. METHODOLOGY

Our methodology comprises several components. First, in Section V-A, we describe extracted MOOC-related features. Next, in Section V-B, we detail the structural features from constructed interaction graph. Finally, in Section V-C, we explain the developed machine learning models for grade prediction.

# A. MOOC-related Features

First, we extracted available MOOC-related features in Table I. Some of these features are specific to a particular entity; for example, the difficulty is specific to the challenge entity. In addition, some features describe the interaction between the entities; for example, the final score is the student's score for a specific challenge.

TABLE I: MOOC-related feat	ares
----------------------------	------

Feature	Description
User ID	Student unique identification number
Challenge $\operatorname{ID}$	Challenge unique identification number
Timestamp	The time when the student first opened the challenge
Final Score	Final score of student on challenge
Exercise ID	Exercise ID for a particular challenge
Course ID	Course ID of a specific exercise
Difficulty	Difficulty level of a challenge
#Retries	Number of attempts to complete a challenge
Duration	Time spent to complete a challenge

# B. Structural Features

As described in Section IV, we constructed a bipartite graph by considering students and challenges as nodes and interactions between these nodes as edges. Table II shows some of the properties of this graph.

From graph BG, we first extracted two node-level pieces of information, including the degree and eigenvector centrality of the challenges and users. The degree of a node is determined by the number of edges incident to it. Eigenvector centrality measures a node's influence within a network [38]. Next, we extracted node embeddings for the users and challenges using two well-known graph representation

TABLE II: Properties of constructed interaction graph

Туре	Bipartite
Includes node features	Yes
Includes edge features	Yes
Number of students $( S )$	5537
Number of challenges $( \mathcal{C} )$	1981
Number of nodes (students+challenges or $ \mathcal{S}  +  \mathcal{C} $ )	7518
Number of interactions (edges or $ E $ )	115124
Density	0.01

learning methods, namely node2vec [22] and DeepWalk [23]. Node2vec is an algorithmic framework for learning continuous feature representations for graph nodes. In node2vec, we learn a mapping of nodes to a low-dimensional space of features that maximizes the likelihood of preserving network neighborhoods of nodes [22]. DeepWalk is a novel approach for learning latent representations of vertices in a network. Statistical models can utilize these latent representations to encode social relations in a continuous vector space [23]. During the experiments, we replaced the user ID and challenge ID in the dataset with these structural features when structural data were to be used. Table III shows the extracted structural features that we employed.

TABLE III: Structural features

Feature	Description
Challenge embedding	Challenge node embedding
User embedding	User node embedding
Degree of challenge	Number of edges linked to a challenge
Degree of user	Number of edges connected to user
Challenge <b>EC</b>	Eigenvector centrality of challenge
User EC	Eigenvector centrality of user

C. Prediction Models

We utilized the following prediction models to predict a student's grade in a challenge:

## Random Forest

Random Forest is a classifier that employs multiple decision trees on various subsets of the input dataset and averages the results to enhance the predicted accuracy. It is based on ensemble learning, which combines multiple classifiers to tackle a challenging problem and improve the model's performance. Formally, given a dataset  $D = \{(x_i, y_i)\}_{i=1}^n$ , Random Forest builds T decision trees  $\{h_t(x)\}_{t=1}^T$  using bootstrapped subsets of D. The final prediction is obtained by averaging the individual tree predictions:

$$H(x) = \frac{1}{T} \sum_{t=1}^{T} h_t(x)$$
(1)

Gradient Boosting

Gradient boosting is a machine learning technique for regression and classification problems that produces a prediction model as an ensemble of weak prediction models. Gradient boosting is an iterative process that effectively transforms a weak learner into a strong learner. Formally, given a dataset  $D = \{(x_i, y_i)\}_{i=1}^n$ , gradient boosting builds an ensemble of weak learners  $\{h_t(x)\}_{t=1}^T$ by minimizing the loss function L(y, F(x)):

$$F(x) = \sum_{t=1}^{T} \alpha_t h_t(x) \tag{2}$$

where  $\alpha_t$  is the step size at iteration t.

XGBoost

XGBoost is an implementation of gradient-boosted decision trees. In this approach, decision trees are generated sequentially. Weights play a crucial role in XGBoost. Each independent variable is weighted before being input into the decision tree that predicts outcomes. Formally, given a dataset  $D = \{(x_i, y_i)\}_{i=1}^n$ , XGBoost builds an ensemble of T decision trees  $\{h_t(x)\}_{t=1}^T$  by minimizing the regularized loss function  $L(y, F(x)) + \Omega(h_t)$ :

$$F(x) = \sum_{t=1}^{T} h_t(x)$$
 (3)

where  $\Omega(h_t)$  is the regularization term that controls the complexity of the decision trees.

**Remarks**. We performed experiments using alternative prediction models, such as Support Vector Machine (SVM) and Decision Tree Classifier (DT), but observed unsatisfactory performance with our data. As a result, we chose to utilize the three prediction models discussed in our analysis. Moreover, our objective is not to develop a novel machine learning methodology. Instead, we aim to examine the influence of structural features in a MOOC-induced graph. Thus, we employed off-the-shelf models. We plan to explore more advanced models in future work.

#### VI. EXPERIMENTS

To evaluate the effectiveness of structural features in challenge grade prediction, we performed various experiments, which are discussed in this section. First, in Section VI-A, we describe the experimental settings. Next, in Section VI-B, we present the main empirical results. Section VI-C includes experimental results showing the association between model predictions across students' academic performance levels. Finally, in Section VI-D, we analyze the importance of different features in grade prediction performance.

### A. Experimental Settings

We divided the final scores into five classes to perform grade classification. The grade range for each class and the number of instances with that grade class is displayed in Table IV.

To split the dataset into train and test sets, for every student, we first sorted the challenges they have completed based on their timestamp. Then, we added approximately the first 80% of challenges for each student in the training set and the rest in the test set. The reason for performing this type of split is to predict a student's future grades based on their previous performance. This way, the prediction reflects the practical usage of grade prediction tasks where we are

TABLE IV: Discretized grades and number of instances in each grade label

Class	Grade range	# Instances
class 0 (Extremely Low)	$0 \leq grade < 20$	34500
class 1 (Very Low)	$20 \leq grade < 40$	34403
class 2 (Average)	$40 \leq grade < 60$	19315
class 3 (Low)	$60 \leq grade < 80$	3970
class 4 (High)	$80 \leq grade \leq 100$	34500

interested in estimating the future grades during an academic period, for example, to intervene and offer help to possible struggling students.

The Random Forest and Gradient boosting methods were implemented using the scikit-learn package. To implement the XGBoost model, we used the XGBoost Python package. In addition, we utilized the node2vec package<sup>1</sup> and the Karateclub package<sup>2</sup> for node2vec and Deepwalk implementation, respectively. The hyperparameters used in learning node embeddings with node2vec and DeepWalk are as follows: embedding dimension: 128, number of walks: 100, walk length: 10, and window size: 10.

**Evaluation Metrics.** The performance prediction models were evaluated using a variety of measures. Suppose TP, FP, TN, and FN are the number of true positive, false positive, true negative, and false negative samples, respectively. Given this notation, we reported the prediction performance using the following metrics.

*Accuracy* calculates the proportion of accurate predictions to the total number of instances considered:

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN}$$

*Precision* is the ratio of accurately predicted positive patterns in a positive class to all predicted positive patterns:

$$Precision = \frac{TP}{TP + FP}$$

*Recall* measures the fraction of positive patterns that are correctly classified:

$$Recall = \frac{TP}{TP + FN}$$

*F1-score* represents the harmonic mean of the Precision and Recall:

$$F1\text{-}score = 2 \times \frac{Precision \times Recall}{Precision + Recall}$$

*ROC* (*Receiver Operating Characteristic*) curve is a graph that displays how well a classification model performs across all classification thresholds. It is a probability curve that plots the TPR (True Positive Rate) and FPR (False Positive Rate) at different threshold values.

*AUC* calculates the total two-dimensional region under the complete ROC curve.

# B. Experimental Results

TABLE V: Comparing the performance of 9 models on grade prediction task. The best algorithm in each column is displayed in **bold**.

Model	Accuracy	Precision	Recall	F1-score
Random Forest	0.80	0.85	0.86	0.86
XGBoost	0.78	0.84	0.85	0.85
Gradient Boosting	0.84	0.89	0.89	0.89
Random Forest + DeepWalk	0.89	0.92	0.92	0.92
XGBoost + DeepWalk	0.87	0.91	0.90	0.90
Gradient Boosting + DeepWalk	0.89	0.92	0.92	0.92
Random Forest + node2vec	0.88	0.92	0.91	0.91
XGBoost + node2vec	0.89	0.92	0.92	0.92
Gradient Boosting + node2vec	0.89	0.92	0.92	0.92

Table V presents the results of nine experiments, as we have three predictive models and three variations of structural features: no structural features, node2vec (+ degree and eigenvector), and DeepWalk (+ degree and eigenvector). From the results presented in Table V, we can make the following observations:

- Incorporating structural features using either DeepWalk or node2vec significantly improves the performance of all three prediction models (Random Forest, XGBoost, and Gradient Boosting) in terms of Accuracy, Precision, Recall, and F1-score.
- Gradient Boosting consistently outperforms Random Forest and XGBoost across all variations, achieving the best overall performance when combined with node2vec structural features.
- There is a close performance between the models incorporating DeepWalk and node2vec features, with only a slight edge in favor of the node2vec-based models.
- ⇒ Given the large number of students enrolled in MOOCs, even a 3% increase in performance is quite significant.
- ⇒ The use of structural features leads to a substantial improvement in the model performance. For instance, the Gradient Boosting model's performance increases from an F1-score of 0.89 to 0.92 when incorporating node2vec features.

These observations highlight the importance of considering complex structural relationships between entities in MOOC grade prediction models and the potential of graph-based representation learning techniques to enhance the performance of these models.

Figure 3 presents the ROC curves and their corresponding AUCs. Since we are addressing a multi-class classification problem, ROC curves can be plotted by comparing one class against the rest. As a result, for each class, we plotted one-versus-rest ROC curves for all nine models and calculated the AUC scores. It can be observed that the predictions for the middle classes are nearly perfect for all models. However, for classes 0 and 4, the models' performance is not as strong. Nevertheless, the ROC curves demonstrate that the models incorporating structural features outperform others in these classes.

<sup>&</sup>lt;sup>1</sup>https://pypi.org/project/node2vec/

<sup>&</sup>lt;sup>2</sup>https://karateclub.readthedocs.io/en/latest/



Fig. 3: One-vs-Rest ROC curves and AUC score for all models.

TABLE VI: Detailed classification results for different Gradient Boosting variations

	Gradi	ient Bo	osting	Gradient	Boostin	g + node2vec	Gradient	Boostin	g + DeepWalk
Class	Precision	Recall	F1-score	Precision	Recall	F1-score	Precision	Recall	F1-score
class 0 (Extremely Low)	0.68	0.72	0.70	0.82	0.72	0.77	0.83	0.72	0.77
class 1 (very low)	0.99	0.99	0.99	0.99	1.00	1.00	0.99	1.00	1.00
class 2 (low)	0.99	1.00	0.99	0.99	1.00	0.99	0.99	1.00	1.00
class 3 (average)	0.99	1.00	1.00	0.99	1.00	1.00	0.99	1.00	1.00
class 4 (high)	0.79	0.74	0.76	0.82	0.88	0.85	0.82	0.89	0.85
accuracy			0.84			0.89			0.89
macro avg	0.89	0.89	0.89	0.92	0.92	0.92	0.92	0.92	0.92
weighted avg	0.84	0.84	0.84	0.89	0.89	0.89	0.89	0.89	0.89

Additionally, we have provided the classification reports (Table VI) and the confusion matrices (Figure 4) for Gradient Boosting (the top-performing model) with and without using structural features to examine their effectiveness on the challenge grade prediction task in greater detail. The primary classification metrics are displayed as a function of each class in the classification report. This offers a more comprehensive understanding of the classifier's behavior than global accuracy, which can conceal functional deficiencies in one class of a multi-class problem. We make the following observations from Table VI and Figure 4:

⇒ We can observe that structural features significantly improve the performance of classes 0 and 4, which the basic

Gradient Boosting model struggles to predict effectively. ⇒ While there has not been a substantial improvement in performance for classes 1, 2, and 3, it's important to note that incorporating structural features makes a notable difference in predicting classes 0 and 4.

- ⇒ The basic Gradient Boosting model faces difficulties in accurately predicting these classes, but the inclusion of structural features has resulted in a 7% improvement in performance for class 0 and a 9% increase for class 4.
- ⇒ It is crucial to emphasize the improvement in class 0, as this class represents students with the lowest grades, whose accurate prediction in automated grade prediction in MOOCs is essential [11, 12]. These students are



Fig. 4: Confusion matrix of (a) Gradient boosting model. (b) Gradient boosting + node2vec model. (c) Gradient boosting + DeepwWalk model.



Fig. 5: Performance of Gradient Boosting and Gradient Boosting + node2vec models for different categories of students.

considered at-risk, and identifying them accurately can enable targeted interventions and support to help them succeed in their courses. The incorporation of structural features in the prediction models proves to be a valuable approach to enhance the prediction performance for this critical group of students. The following section offers an in-depth explanation of our automated challenge grade prediction system across student groups with different academic attainment levels.

#### C. Grade prediction and student academic performance

As discussed in Section I, the primary goal of grade prediction is typically to identify and support students with low academic performance. It is essential for a grade prediction model to accurately predict future grades for struggling students. To evaluate the performance of our best model (Gradient Boosting) for low-performing students, we categorized students into five groups based on their grades in the training set, which can be considered as their previous performance. Students with more than 90% of their grades in classes 0 and 1 (low grades) were labeled as "Extremely Low" performing. Similarly, students with 80-90% low grades were labeled "Very Low", 50-80% low grades as "Low", 20-50% low grades as "Average", and students with less than 20% low grades were labeled "High" performing.

Figure 5 illustrates the performance of Gradient Boosting with and without using structural features for the five different categories of students. In terms of Weighted F1-score, the prediction for "Extremely Low" performing and "Very Low" performing students improved by 6% and 4%, respectively, when structural features were included. This, again, demonstrates the positive impact of incorporating structural features for predicting the performance of struggling students.

## D. Feature Analysis

Figure 6 shows the feature importance in different Gradient Boosting models. In the first model, challenge ID, duration, and timestamp are the top three factors in predicting the challenge grade. However, in the second model, the embedding of challenge nodes along with user node embeddings and degree of challenge node dominates the feature importance plot. In other words, structural features become prominent features that surpass ordinary MOOC-related features, e.g., challenge ID. Finally, in the third model, where DeepWalk was used to obtain the embeddings, importance is more distributed but still the challenge embeddings, and user embeddings along with node degree and eigenvector centrality are critical factors in the model's predictions. This, again, emphasizes the importance of considering structural features in grade prediction.

#### VII. CONCLUSION

MOOC is one of the most exciting recent phenomena in higher education. The promise of MOOC is the democratization of education [2]. It has attracted millions of learners worldwide. The high enrollment, diversity of topics, and the scarcity or lack of instructors and student-instructor interaction have automated some tasks. In particular, grade prediction is a crucial task providing several advantages. However, in this study, we identified two issues associated with grade



Fig. 6: Feature importance of models (a) Gradient boosting model. (b) Gradient boosting + node2vec model. (c) Gradient boosting + DeepWalk model.

prediction in MOOC: 1) the lack of enough granularity in the prediction (i.e., focusing on high-level tasks such as course prediction), and 2) ignoring complex yet effective structural relationships among pertinent entities. In this study, we first introduced a novel and large MOOC dataset providing multiple aspects of academic data to overcome these issues. Notably, challenges are short and small exercises focusing on a specific skill. Since this is a complicated problem and offers high granularity in the prediction, we pressed onward to predict students' challenge grades. We extracted salient and dense structural representations from a constructed interaction graph between students and challenges to address the second issue and improve prediction performance. We performed extensive experiments and showed that structural features could significantly improve the quality of the challenge prediction.

There are several intriguing future directions. First, an interesting research question is how our challenge prediction system is perceived in real-world settings. In fact, our research group is currently investigating this aspect. Second, one could extend the interaction graph beyond a bipartite graph and include other entities, such as courses and disciplines. Third, such a heterogeneous graph could be used to perform other tasks, like course recommendation. Finally, exploring more advanced machine learning methods, such as node representation learning in heterogeneous graphs [39], applied to the graph in our dataset, is a promising direction for future research.

#### ACKNOWLEDGMENT

In this work, Dr. Jiangtao Huang is supported by the National Natural Science Foundation of China (62067007).

#### References

 G. Christensen, A. Steinmetz, B. Alcorn, A. Bennett, D. Woods, and E. Emanuel, "The mooc phenomenon: Who takes massive open online courses and why?" *Available at SSRN 2350964*, 2013.

- [2] T. Dillahunt, Z. Wang, and S. D. Teasley, "Democratizing higher education: Exploring mooc use among those who cannot afford a formal education," *International Review* of Research in Open and Distributed Learning, vol. 15, no. 5, pp. 177–196, 2014.
- [3] E. Ossiannilsson, "Moocs for lifelong learning, equity, and liberation," *MOOC (Massive Open Online Courses)*, 2021.
- [4] S. Papadakis, "Moocs 2012-2022: An overview," Advances in Mobile Learning Educational Research, vol. 3, no. 1, pp. 682–693, 2023.
- [5] N. M. Castillo, J. Lee, F. T. Zahra, and D. A. Wagner, "Moocs for development: Trends, challenges, and opportunities," *International Technologies & International Development*, vol. 11, no. 2, p. 35, 2015.
- [6] T. Daradoumis, R. Bassi, F. Xhafa, and S. Caballé, "A review on massive e-learning (mooc) design, delivery and assessment," in 2013 eighth international conference on P2P, parallel, grid, cloud and internet computing. IEEE, 2013, pp. 208–213.
- [7] A. McAuley, B. Stewart, G. Siemens, and D. Cormier, "The mooc model for digital practice," 2010.
- [8] L. Czerniewicz, A. Deacon, J. Small, and S. Walji, "Developing world moocs: A curriculum view of the mooc landscape," 2014.
- [9] Q. Qi, Y. Liu, F. Wu, X. Yan, and N. Wu, "Temporal models for personalized grade prediction in massive open online courses," in *Proceedings of ACM Turing Celebration Conference-China*, 2018, pp. 67–72.
- [10] J. Fan, Y. Jiang, Y. Liu, and Y. Zhou, "Interpretable mooc recommendation: a multi-attention network for personalized learning behavior analysis," *Internet Research*, vol. 32, no. 2, pp. 588–605, 2022.
- [11] T.-Y. Yang, C. G. Brinton, C. Joe-Wong, and M. Chiang, "Behavior-based grade prediction for moocs via time series neural networks," *IEEE Journal of Selected Topics* in Signal Processing, vol. 11, no. 5, pp. 716–728, 2017.
- [12] M. Adnan, A. Habib, J. Ashraf, S. Mussadiq, A. A. Raza, M. Abid, M. Bashir, and S. U. Khan, "Predicting atrisk students at different percentages of course length for

early intervention using machine learning models," *Ieee Access*, vol. 9, pp. 7519–7539, 2021.

- [13] D. J. Lemay and T. Doleck, "Grade prediction of weekly assignments in moocs: mining video-viewing behavior," *Education and Information Technologies*, vol. 25, pp. 1333–1342, 2020.
- [14] P. G. de Barba, G. E. Kennedy, and M. D. Ainley, "The role of students' motivation and participation in predicting performance in a mooc," *Journal of Computer Assisted Learning*, vol. 32, no. 3, pp. 218–231, 2016.
- [15] B. Xu and D. Yang, "Motivation classification and grade prediction for moocs learners," *Computational intelli*gence and neuroscience, vol. 2016, pp. 4–4, 2016.
- [16] Z. Ren, H. Rangwala, and A. Johri, "Predicting performance on mooc assessments using multi-regression models," arXiv preprint arXiv:1605.02269, 2016.
- [17] Y. Chen and M. Zhang, "Mooc student dropout: Pattern and prevention," in *Proceedings of the ACM Turing 50th Celebration Conference-China*, 2017, pp. 1–6.
- [18] W. Xing and D. Du, "Dropout prediction in moocs: Using deep learning for personalized intervention," *Journal of Educational Computing Research*, vol. 57, no. 3, pp. 547–570, 2019.
- [19] J.-W. Tzeng, C.-A. Lee, N.-F. Huang, H.-H. Huang, and C.-F. Lai, "Mooc evaluation system based on deep learning," *International Review of Research in Open and Distributed Learning*, vol. 23, no. 1, pp. 21–40, 2022.
- [20] A. Chauhan, "Massive open online courses (moocs): Emerging trends in assessment and accreditation," *Digital Education Review*, no. 25, pp. 7–17, 2014.
- [21] P. M. Moreno-Marcos, C. Alario-Hoyos, P. J. Muñoz-Merino, and C. D. Kloos, "Prediction in moocs: A review and future research directions," *IEEE transactions on Learning Technologies*, vol. 12, no. 3, pp. 384–401, 2018.
- [22] A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016, pp. 855–864.
- [23] B. Perozzi, R. Al-Rfou, and S. Skiena, "Deepwalk: Online learning of social representations," in *Proceedings* of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining, 2014, pp. 701– 710.
- [24] H. Karimi, J. Huang, and T. Derr, "A deep model for predicting online course performance," *Association for the Advancement of Artificial Intelligence*, 2020.
- [25] H. Karimi, T. Derr, J. Huang, and J. Tang, "Online academic course performance prediction using relational graph convolutional neural network." *International Educational Data Mining Society*, 2020.
- [26] H. Karimi, T. Derr, K. T. Torphy, K. A. Frank, and J. Tang, "A roadmap for incorporating online social media in educational research," *Teachers College Record*, vol. 121, no. 14, pp. 1–24, 2019.
- [27] H. Karimi, T. Derr, K. Torphy, K. A. Frank, and J. Tang, "Towards improving sample representativeness of teach-

ers on online social media: A case study on pinterest." Artificial Intelligence in Education: 21st International Conference, AIED 2020, Ifrane, Morocco, July 6–10, 2020, Proceedings, Part II 21, 2020, pp. 130–134.

- [28] H. Karimi, K. T. Knake, and K. A. Frank, "Teachers in social media: A gender-aware behavior analysis," in 2022 IEEE International Conference on Big Data (Big Data). IEEE, 2022, pp. 1842–1849.
- [29] K. T. Knake, H. Karimi, S. Hu, K. A. Frank, and J. Tang, "Educational research in the twenty-first century: Leveraging big data to explore teachers' professional behavior and educational resources accessed within pinterest," *The Elementary School Journal*, vol. 122, no. 1, pp. 86–111, 2021.
- [30] H. Karimi, J. Tang, X. Weiss, and J. Huang, "Automatic identification of teachers in social media using positive unlabeled learning," in 2021 IEEE International Conference on Big Data (Big Data). IEEE, 2021, pp. 643–652.
- [31] A. Ramesh, D. Goldwasser, B. Huang, H. Daumé III, and L. Getoor, "Modeling learner engagement in moocs using probabilistic soft logic," in *NIPS workshop on data driven education*, vol. 21, 2013, p. 62.
- [32] C. G. Brinton and M. Chiang, "Mooc performance prediction via clickstream data and social learning networks," in 2015 IEEE conference on computer communications (INFOCOM). IEEE, 2015, pp. 2299–2307.
- [33] C. Piech, J. Bassen, J. Huang, S. Ganguli, M. Sahami, L. J. Guibas, and J. Sohl-Dickstein, "Deep knowledge tracing," Advances in neural information processing systems, vol. 28, 2015.
- [34] B.-H. Kim, E. Vizitei, and V. Ganapathi, "Gritnet 2: Real-time student performance prediction with domain adaptation," *arXiv preprint arXiv:1809.06686*, pp. 1–8, 2018.
- [35] H. Li, H. Wei, Y. Wang, Y. Song, and H. Qu, "Peerinspired student performance prediction in interactive online question pools with graph neural network," in *Proceedings of the 29th ACM International Conference* on Information & Knowledge Management, 2020, pp. 2589–2596.
- [36] J. Fang, L. Tang, J. Yang, and M. Peng, "Social interaction in moocs: The mediating effects of immersive experience and psychological needs satisfaction," *Telematics and Informatics*, vol. 39, pp. 75–91, 2019.
- [37] S. Shatnawi, M. M. Gaber, and M. Cocea, "Automatic content related feedback for moocs based on course domain ontology," in *Intelligent Data Engineering and Automated Learning–IDEAL 2014: 15th International Conference, Salamanca, Spain, September 10-12, 2014. Proceedings 15.* Springer, 2014, pp. 27–35.
- [38] P. Bonacich, "Some unique properties of eigenvector centrality," *Social networks*, vol. 29, no. 4, pp. 555–564, 2007.
- [39] X. Wang, D. Bo, C. Shi, S. Fan, Y. Ye, and S. Y. Philip, "A survey on heterogeneous graph embedding: methods, techniques, applications and sources," *IEEE Transactions* on Big Data, 2022.