Paper ID #51191

A Time Evolving Link Prediction Framework in Social Learning Networks for In-Person Courses

Jeremy Wong, University of California, San Diego

Computer Engineering master's student at University of California San Diego with experience in machine learning and neural networks. Working as a research assistant under Dr. Rajeev Sahay to understand how machine learning can be used to improve academic outcomes.

Ryan Joseph Irwandy, University of California, San Diego
Dr. Rajeev Sahay, University of California, San Diego

Rajeev Sahay received the B.S. degree in electrical engineering from The University of Utah, Salt Lake City, UT, USA, in 2018, and the M.S. and Ph.D. degrees in electrical and computer engineering from Purdue University, West Lafayette, IN, USA, in 2021 and 2022, respectively. Currently, he is a faculty member in the Department of Electrical and Computer Engineering at UC San Diego. In 2021, he was the recipient of the Purdue Engineering Dean's Teaching Fellowship. Dr. Sahay's teaching interests include introductory and advanced programming (in C, C++, and Python) in the electrical engineering curriculum. His research interests lie in the intersection of networking and machine learning, especially in their applications to wireless communications and engineering education.

A Time Evolving Link Prediction Framework in Social Learning Networks for In-Person Courses

Abstract

Collaborative learning, where students work with peers to exchange ideas and solve problems, is widely recognized as a driver of academic success. However, while instructors routinely promote collaboration, there is little methodology for systematically modeling in-person peer interactions. Social learning networks (SLNs) offer a computational framework for representing collaboration as graphs, with students as nodes and their interactions as weighted edges. However, existing SLN research has focused almost exclusively on virtual environments, such as massive open online courses (MOOCs) and learning management system (LMS) forums, where digital traces make interactions easy to capture and analyze. As a result, link prediction, a common network analysis task used to forecast future connections, has never been applied to a physical classroom, where face-to-face interactions are not automatically recorded and quantified. This work introduces the first framework for constructing an SLN in a physical classroom and for performing link prediction on in-person student interactions. To model links without digital logs, we collect weekly student evaluations of their collaborative learning exercises and use these self-reports to build weighted graphs that reflect both the frequency and quality of peer collaboration over time. Using this SLN, we investigate how behavioral links relate to academic outcomes and evaluate whether link prediction can help form more effective student pairs. We implement a graph neural network (GNN) approach using a modified weighted GraphSAGE architecture and apply it to data from a real electrical engineering course at an American R1 institution. Our method achieves a link-prediction AUC of 0.8000, outperforming multiple state-of-the-art baselines. These results demonstrate that in-person interactions can be meaningfully modeled as an SLN and leveraged to provide instructors with data-driven recommendations for beneficial student pairings.

1 Introduction

Numerous studies have been devoted to developing a greater understanding of the factors that lead to student academic success and how these factors can be used to predict this success. These studies aim to help educators intervene for students who struggle in their courses early on to provide students with better chances of success at the end of the course [1]. One particular approach used by instructors to improve course outcomes is peer-to-peer interactions, which are known to positively impact student learning [2]. In this work, we focus on analyzing, forming, and predicting such meaningful interactions and connections between students in an attempt to promote positive collaborative learning experiences and thus improve academic outcomes.

One of the main challenges with predicting student performance is the type of data typically used for academic performance analysis. Commonly analyzed metrics and predictors of a student's academic performance consist of standardized examinations, admissions scores, and previous coursework performance [1], but these metrics are redundant, as students who perform well in earlier academic settings are likely to continue to perform well in more advanced academic courses and environments later on. Beyond academic metrics, the background, personality, behavioral patterns, and social activities of a student also affect the student's learning outcomes [3] [4]. Social activities are the least quantifiable of these influences, which can make predicting student success more challenging with respect to this variable. However, many have studied the benefits of cooperative peer-to-peer interactions in improving student academic success [5], and with the ever-growing popularity of online forums and the greater reliance on learning management systems (LMSs), researchers have used metrics collected from these websites to analyze the effects of interactions through these media on student learning [6] [7]. Most current research is limited to these online interactions, where data is often captured automatically and readily available for analysis. Contrary to online courses, in-person classrooms do not have this data, both due to the difficulty in quantifying physical interactions and the lack of automated activity tracking. Despite these challenges, understanding how student behavior and interactions inform future collaborative behaviors and ultimately affect academic performance in in-person classrooms is important as (i) physical behavior can differ greatly from online behavior and (ii) the vast majority of higher education institutions deliver instruction primarily via in-person modalities with only supplemental online components.

We address the challenges of capturing and analyzing in-person classroom interaction data by establishing a novel method for capturing physical behavioral patterns. In doing so, behavioral data can be gathered from courses with in-person modalities rather than from online platforms and can be used exclusively for the analysis of student performance. Thus, specifically for an in-person classroom, we focus on the following research question: *How does student behavior in in-person courses predict academic outcomes?* and, as a follow-up to that question, *How can effective student collaborations in in-person courses be accurately quantified to inform optimal collaborative learning experiences?* The former question is grounded in using social behavior, rather than academic measures, to analyze student performance by analyzing behavioral data collected from lab groups in a physical classroom. The latter question paves the path for designing a tool to optimize the formation of these groups for the best overall learning outcomes, since random assignment of groups may not benefit or may even harm student learning.

To this end, we consider using *Social Learning Networks* (SLNs), graphs in which nodes represent students and edges (i.e., *links*) represent student interaction, to represent classroom interactions. In online classrooms, where student interaction data is automatically logged via digital interactions (e.g., on discussion forums), SLNs are formed naturally. However, implementing SLNs in in-person courses is more challenging compared to digital settings because there is often no measure of physical interactions in in-person courses. As a result, the task of link prediction, which aims to compute the probability of unobserved or future interactions between students based on their prior social interactions, in in-person courses via SLNs has typically been considered unviable. In this work, we develop a framework for measuring both the existence and the quality of student interactions, which, in turn, are used to perform effective link prediction in a physical classroom, thus leading to optimized collaborative learning experiences.

Overall, we make the following contributions in this work:

- To the best of our knowledge, we, for the first time, apply a data-driven social learning network to an in-person classroom to model in-person interactions between students as opposed to online interactions via a digital forum.
- We develop an academic performance prediction framework based on self-reported behavioral data of students to provide an early intervention method for students who may be struggling in a given course.
- We perform an empirical evaluation of our proposed framework on a real-world electrical engineering classroom.

2 Related Work

This section explores the current literature pertaining to the focus of this study. First, Sec. 2.1 discusses how student behavioral patterns have been observed to have an effect on academic performance. Next, Sec. 2.2 describes studies on how academic outcomes are influenced by student collaboration and social engagement in their courses and institutions. Lastly, Sec. 2.3 explores social learning networks as a model for student-to-student interactions in classrooms and how SLNs have historically been used to predict future interactions in digital settings.

2.1 Behavioral Data-Based Academic Performance

Student behavioral patterns have been shown to have a large influence on academic performance. Several studies focus on gathering data from online classrooms or systems, where metrics can be easily collected. [8] uses time spent on the LMS to quantify engagement, finding that increased student engagement with the application predicted better academic outcomes. However, this metric does not factor in the quality of the course interaction, only the duration. The data from [9] was collected from LMS data in a flipped classroom, finding that increased engagement with certain academic activities improved course performance. More research explores the use of machine learning (ML) or deep learning (DL) models for analyzing virtual learning environments [10] [11] [12]. The data includes demographic and academic features, as well as features centered on interaction and engagement with course learning platforms, of which different combinations of these subsets are evaluated. Others use data such as grade data, written responses, attendance, or even time spent on the LMS to train and test their model on [13]. These works, however, partially rely on academic data to predict academic success, adding redundancy to their models. Additionally, these studies focus solely on online platforms, where student behavior can differ greatly from in-person classrooms. These papers also hone in on student behavioral patterns individually as they interact with their academic coursework or online material, not how they interact with one another.

2.2 Collaborative Learning

Collaborative learning emphasizes the importance of peer-to-peer interactions. Previous works have revealed collaborative learning involves students more actively in the learning process, fosters positive student-to-student and student-to-teacher relationships, and improves classroom academic results [14] [15]. Students who are motivated to learn and work together are more likely to succeed in class [16]. Researchers in multiple studies sought to discover how metrics obtained from LMSs can be leveraged to understand how academic engagement and student social interactions through the medium of an LMS affect academic performance. In [17], each student self-reported academic participation in written reflections, in which each of the following four categories of engagement was extracted using natural language processing: behavioral, emotional, social, and cognitive. In other words, the social engagement of the students in the classroom was measured using the written responses of the students. However, this method can only estimate the quality of student interactions, not directly measure it, and uses the same data point to extract all four subcategories of academic engagement instead of separate data points for each category. Another study explores the impact of student academic involvement on student social relationships and academic performance [18]. Using Facebook as the social network, the study faced difficulties with conflicting predictors, as collaborative academic interactions between students positively predicted academic success, while more time spent on the platform negatively predicted academic success. Thus, using Facebook as a social network, not specifically a social learning network, blurs the focus on academic interactions. [19] discussed an academic performance predictor that used friendship and peer interaction data, noting that students who worked well with their peers and those who had friends of certain backgrounds generally performed better in their courses. One shortcoming of this study was that the predictor also used demographic data, which is often inaccessible. All these studies reveal that increased engagement with other students generally improves academic performance, but few propose how peer-to-peer interactions can be created or enhanced. However, many works have been dedicated to social learning networks that model student interactions in academic contexts.

2.3 Social Learning Networks

Social learning networks have been the focus of several studies that seek to understand the interactions of students in massive open online courses (MOOCs) and the effects of these interactions on grade outcomes [7] [20]. However, these papers employ models designed for large amounts of data due to their large datasets; analyze online social learning networks, not in-person classroom settings; and rely on grade data to perform their analysis, which may not be available due to privacy concerns. Despite these differences, previous research has demonstrated the ability to model SLNs in a graph structure with nodes and edges, on which link prediction, which has been extensively applied to social networks such as social media platforms [21], can be performed. The node features and link values of weighted and directed graphs can be used as input to machine learning models to learn and form predictions on complex network structures [22] [23] [24]. In the context of education, however, these problems are neither designed for smaller classrooms nor explore graphs without node features, both of which are important for courses where enrollments are low, and student data must be kept private. SLNs have been used to analyze student interactions, although none explore these networks in physical academic settings due to the difficulty of obtaining data on or quantifying in-person student behavior, let alone analyzing and performing link prediction on such data [25]. In this work, we address these challenges by developing a link prediction SLN framework, using purely behavioral data, which can operate on small courses and accounts for the issue of training the model with a cold start due to the lack of previous data.

3 Methodology

This section describes the proposed framework in this study. Sec. 3.1 starts by discussing the data collection process in this work. Next, Sec. 3.2 explains how a graph is used to represent the behavioral data. Sec. 3.3 provides an overview of the time evolution of the social learning network as more links are formed. Following this description, our approach of using graph neural networks to perform link prediction is explained in Sec. 3.4. Sec. 3.5 closes with a discussion of the training regime and the metrics used to evaluate the performance of our framework.

3.1 Data Format

The data for the experiments was gathered through self-reported surveys. These surveys were administered for each collaborative assignment, with each assignment due on a weekly basis. Each week, the students were randomly assigned a new partner with whom to collaborate. After each session of working together, the students were asked to rate their experiences with their partners on a scale from 1 to 10 based on the following questions:

- *To what extent did the interaction help you understand the course material better?*
- *How engaged did you feel during the interaction?*
- *How equally do you think you both contributed to the discussion?*
- *How well were you able to communicate your ideas during the interaction?*
- *How would you rate the flow of the conversation?*
- *How would you rate your overall satisfaction in working with your assigned partner today?*

Each student was also assigned a lab number that served as a personal identification number for the course, allowing all personally identifiable information to be removed from the data. This number, along with the lab number of the student's partner, was also recorded in the survey. By collecting the data in this manner, the data could then be efficiently represented using a graph, with the student and their partner(s) represented by nodes, and the numeric ratings of each student represented by links between the nodes.

3.2 Graph Representation of Data

We use the responses to the surveys to represent a weighted directed graph, where students are the nodes and their survey responses are weighted edges. Because the survey responses are non-binary, edges are weighted, and since each student completed an evaluation of their partner, edges are directed. Mathematically, this graph $\mathcal{G}$ has N nodes, where N is the number of students in the course. Thus, the set of students $\mathcal{S}$ can be notated as $\{S_1, S_2, ..., S_N\}$. An edge is notated as

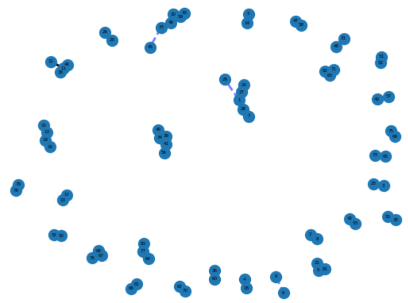

(a) SLN earlier in course

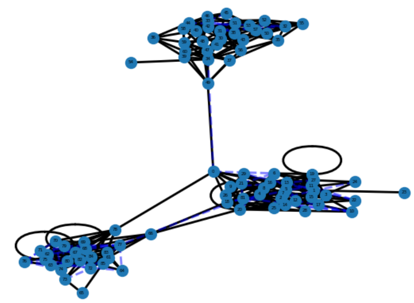

(b) SLN later in course

Figure 1: SLN graphs from earlier and later in the course. The earlier graph visualizes the beginning of the course when students have only interacted with one or two other students. The later graph shows the end of the course, when students have interacted with several other students in one of three distinct clusters.

(u, v), S_u representing the student who completed the survey, and S_v their lab partner. Thus, $1 \leq u, v \leq N$ and $u \neq v$, so there are $N(N-1)$ permutations of node pairs. These node pairs are assigned weights $w_{u,v}$, where non-zero weights represent pairs of students who have interacted and zero weights represent pairs of students who have not. Each weight $w_{u,v}$ in the set of weights $\mathcal{W}$ are also normalized to be in $[0.1, 1.0]$ by dividing each student's response by 10. For example, if student S_1 completes the survey with student S_2 as their partner and gives a score of 8 for one of the questions, the pair would be represented as $(1, 2)$ and the weight as $w_{1,2} = 0.8$. Likewise, if student S_2 gives a score of 6 for one of the questions to student S_1, the pair would be represented as $(2, 1)$ and the weight as $w_{2,1} = 0.6$.

3.3 Link Formation Over Time

To understand the time evolution of the social learning network explored in this study, Fig. 1 visualizes the SLN at the beginning and near the end of the course. Dotted lines represent a "negative" student interaction, or an interaction scored between week 1 through week 5 of the course, while solid lines represent a "positive" student interaction, or an interaction scored between week 6 through week 10 of the course. The SLN at the beginning of the course shows the graph in an early stage, where few edges are formed between pairs of students. However, by the end of the course, the graph reveals three distinct clusters of students, corresponding to the course having three labs sections. It can be seen that student interactions across different sections were rare due to the random pair assignments occurring within a single lab section. Understanding the dynamic nature of the SLN in this study provides a foundation for the temporal analysis of the link prediction models.

Our SLN model also included "negative links," which were zero-valued links to help inform better link prediction by balancing the classes. Especially in smaller classrooms with fewer data points, balancing the frequency of classes that represent positive and negative student interactions is important to improve model prediction. Self-reported data in in-person classrooms also tends to be messy and skewed toward one end of the scale, which is another reason why using negative

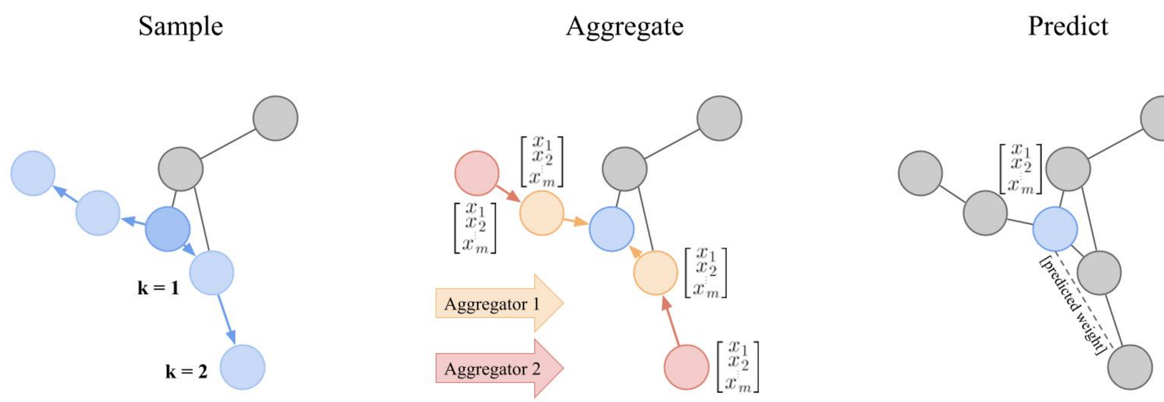

Figure 2: Visualization of sample neighborhood, neighbor feature information aggregation, and graph context and label prediction steps of the GraphSAGE algorithm [26].

links is beneficial for this application.

3.4 Graph Link Prediction

With the data now in graph form, it can be input into the link prediction model. Here, we consider a graph neural network (GNN) for link prediction.

GNNs are an embedding-based and deep learning-based method designed to train on graph nodes and links [27]. This study employs the GraphSAGE algorithm, which is a framework that utilizes feature information from nodes, or students, to form node embeddings for hidden data [26]. GraphSAGE does not train embeddings for each node individually, but instead learns a function that creates embeddings through the sampling and aggregation of features from the local neighborhood of a given node. The steps in this process are visualized in Fig. 2. As the algorithm iterates, nodes aggregate more information from their direct neighbors, resulting in any given node learning more about the entire graph structure. In the context of an SLN, the student nodes gradually share and receive more knowledge about other student nodes and their connections or interactions. Naturally, based on the discussion in Sec. 3.3, the GNN is expected to perform poorly in the earlier stages of the graph due to its cold start, meaning it has no previous data on which to train, but should perform better as time progresses (i.e., in later stages of the course). Modified from the standard GraphSAGE, this model accepts weighted edges instead of only binary edges, allowing the range of scores 0.1 to 1 to be used as input to the model. The following steps comprise a single layer in the modified GraphSAGE model.

3.4.1 Pre-transformation of Node Features

Given input node features $x \in \mathbb{R}^{N \times F}$, where N is the number of nodes, F is the number of features, and F' is the number of features after transformation, to determine the node embeddings, for each student $\mathcal{S}_u$, a learned linear projection of the input features is computed. This linear projection is expressed by

$$\tilde{x}_u = W x_u + b \tag{1}$$

where $W \in \mathbb{R}^{F' \times F}$, and $\tilde{x}_u \in \mathbb{R}^{F'}$ are the transformed features and $b \in \mathbb{R}^{F}$ is the bias term in the linear projection used in the message passing phase.

3.4.2 Message Passing

The message passing phase can be divided into three parts: computing the message for each edge, aggregating the messages for each node, and updating each node.

For each edge, or student interaction, (u, v), a message is computed by the product

$$m_{uv} = w_{uv} \cdot \tilde{x}_v, \tag{2}$$

which scales each neighbor, or partner, $\mathcal{S}_j$'s message by the edge weight or the interaction score w_{uv}. x_v represents the partner's transformed feature.

After computing the message for each student interaction, for each student S_u, messages from all partners $v \in \mathcal{N}(u)$ are aggregated by summation:

$$m_u = \sum_{v \in \mathcal{N}(u)} m_{uv} = \sum_{v \in \mathcal{N}(u)} w_{uv} \cdot \tilde{x}_v \tag{3}$$

Finally, the update step concatenates the original node features x_u with the aggregated message m_u, then applies another linear transformation:

$$x'_u = W_{update} \cdot [x_u \parallel m_u] + b_{update} \tag{4}$$

where $[x_u \parallel m_u] \in \mathbb{R}^{F+F'}$ is the concatenation and $W_{update} \in \mathbb{R}^{F' \times (F+F')}$. This result serves as a learned combination of the node's own identity and the partner neighborhood information.

3.5 Training and Evaluation

With the model implemented, it can now be trained and evaluated on the data.

3.5.1 Training

For training, the GNN used the Adam optimizer [28], which is one of the most commonly used deep learning optimization methods due to its efficiency and performance in various applications [29]. This paper also employs Mean Squared Error (MSE) Loss for the loss function. The GNN was also trained and evaluated twice: once with raw input data and once with binarized input data, where scores 0 through 5 were set at 0 and scores 6 through 10 were set to 1. These will be described as "GNN" and "GNN (bin)," respectively, in 1.

3.5.2 Evaluation

To evaluate the performance of the models, this study measures both accuracy and the area under the curve metrics, and how they change over time for each model. The following acronyms will be used to describe these metrics: true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN).

Accuracy: This metric compares the predicted outputs with the ground truth outputs. It is defined as the number of correctly predicted outputs divided by the total number of predictions. This can be expressed as

$$Acc = \frac{TP + TN}{TP + TN + FP + FN} \tag{5}$$

Area Under the Curve (AUC): This metric is defined as the area under the receiver operating characteristic (ROC) curve. True positive rate (TPR) and false positive rate (FPR) can be computed using the following expressions:

$$FPR = \frac{FP}{TN + FP} \qquad TPR = \frac{TP}{TP + FN} \tag{6}$$

In the case of random selection, the AUC should be 0.5, while a perfect classifier will achieve an AUC of 1. Most classifiers will have $0.5 < AUC < 1$, which means they have some ability to classify correctly. AUC is more helpful than accuracy when there is a large imbalance between classes in a dataset because it provides a measurement of how well a model can distinguish between classes rather than how often the model is correct. SLNs naturally have these large imbalances due to their sparsity by design, as there are significantly more unformed links than formed links between students. The advantage of using AUC to evaluate the model performance will be further highlighted in Sec. 4.

Finally, we examine the progression of accuracy and AUC over time. Because the data was collected periodically, the SLN adds new connections each period. By splitting all of the collection periods into training and testing sets, the effect of more data, which corresponds to a longer link formation time, can be explored. The train-test split iterates from one period of training data and the remaining periods of testing data to all but one period of training data and the last period of testing data, increasing the training set and decreasing the testing set by one period each iteration. For each train-test split, the accuracy and AUC are computed and compared.

4 Results and Discussion

In this section, the empirical analysis of our framework discussed in Sec. 3 is performed. Specifically, Sec. 4.1 describes the setup for the empirical analysis, detailing the data collection, model architecture, and training regime. Next, Sec. 4.2 discusses the relationship found between behavioral patterns and academic performance. Sec. 4.3 outlines the baseline models used to benchmark the performance of our model. Lastly, Sec. 4.4 describes the performance of the main framework in comparison to the baselines and highlights some key results.

4.1 Empirical Setup

As mentioned previously, the data used in this study was collected through student self-reporting, specifically through a Google Form survey. This survey was administered after weekly lab sections for nine weeks in a ten-week lower-division electrical engineering course at the

University of California San Diego, focusing on computer engineering and computer architecture fundamentals. The course had an enrollment of 86 students, which is small relative to many other courses, especially online courses that can have hundreds or thousands of students. The students in this course ranged from first-year to fourth-year, and were divided into three lab sections and assigned random partners within their designated sections. During each weekly section, the students were tasked with completing an assignment with their partners. No guidelines or policies were enforced for how the students were to collaborate. After completing the assignment (or at the end of the lab section), students would complete the Google Form survey to evaluate their experiences for that week. The responses collected through these surveys were used to quantify students collaborative satisfaction in working with their partners and form the SLN.

The GraphSAGE GNN model used to analyze the SLN and perform link prediction consisted of two layers. The complete model is formed by the first message passing layer, which receives the input, a ReLU activation layer, and the second message passing layer, which computes the output. The model uses 16 hidden layers, and the Adam optimizer utilizes a learning rate α of 0.01. The model is trained for 50 epochs using the full dataset for each training batch. This training regime iterates eight times, using more data for each subsequent iteration. Initially, the model is trained using only the first week of data, while the remaining eight weeks of data are used in testing. The next iteration uses the first two weeks of data in the training set and the remaining seven weeks in the testing set. This process repeats until the first eight weeks of data are included in the training set and the last week in the testing set.

4.2 Behavioral Pattern-Based Academic Performance

Here, we focus on our first research question and analyze how student behavior predicts academic outcomes. To analyze this effect, we correlate the week-to-week behavioral results that the six different dependent variables (i.e., the answers to the survey questions shown in Sec. 3.1) have on a student's final grade in the course. We use principal component analysis (PCA) to project the six variables to a 2-D plane for visualization. The findings of this analysis raise questions about possible uses of behavioral data: *Can students' attitudes or behaviors toward a course be predicted? Can course performance prediction be performed using behavioral data? Can it be predicted with whom students are most likely to work well?* Using the same behavioral data as the data in this analysis, we address the first of these two questions in the subsequent paragraph and the last questions in the subsequent subsections. Notably, we do not use grade data to predict academic success due to the small class size and the abundant literature on grade prediction using prior grade data.

Fig. 3 shows the weekly visualizations of the behavioral data (after using PCA to project the data onto two dimensions). Each point, representing a student, was given a color that corresponded to the student's cumulative lab grade up to that week. *Note: the number of participants in the data changes from week to week due to absences, students dropping the class, or late enrollments.*

Through the plots in Fig. 3, we see that there are fairly well-defined clusters which demonstrate the grade similarity of students who had similar qualities of collaborative experiences. Hence, there is a non-negligible effect of behavior on student academic outcomes, which verifies the

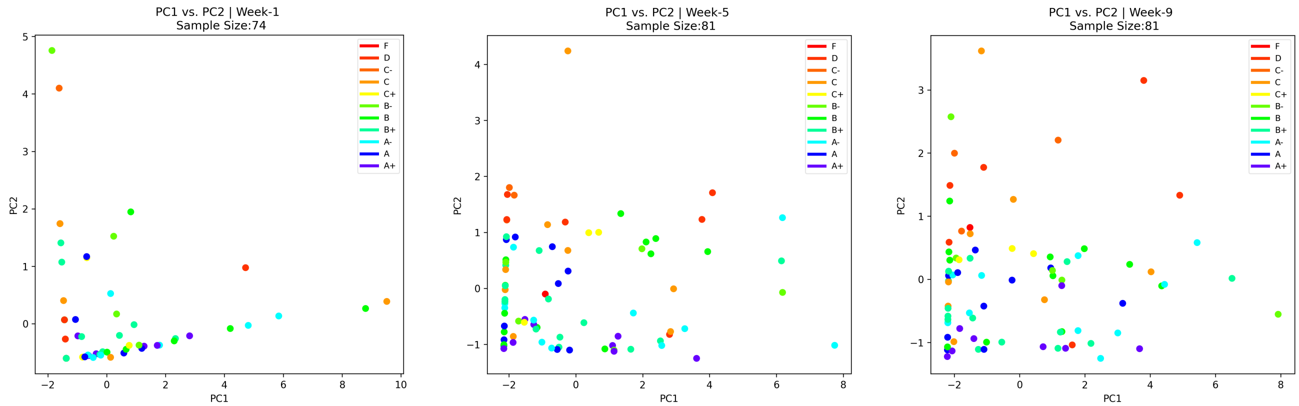

(a) PCA plot for Week 1 (b) PCA plot for Week 5 (c) PCA plot for Week 9

Figure 3: PCA plots for weeks 1, 5, and 9. The more blue a point is, the higher the grade (A+, A, A-, B+, etc.), and the more red a point is, the lower the grade (C, C-, D, etc.). For weeks 1 through 9, the trends of the principal components from the training and testing datasets were generally the same. That is, higher grades tended to cluster toward the bottom of the graph, while the opposite was true for lower grades, demonstrating a relationship between behavior and academic success.

correlation between behavioral data and student learning and also showcases this relationship using our proposed behavioral data collection methodology. With this confirmation from the data collected from the in-person classroom, the first research question, *How does student behavior in in-person courses predict academic outcomes?* is answered, and the method for quantifying the quality of physical interactions between students is shown to be effective. Driven by the answer to the first, the second question, *How can effective student collaborations in in-person courses be accurately quantified to inform optimal collaborative learning experiences?* will be addressed in subsequent sections.

4.3 Baseline Models

Here, we outline the baseline models considered as a point of comparison for our proposed GNN-based link prediction framework.

Convolutional Neural Network (CNN): This model, composed of convolutional filter layers that downsample and upsample as needed [30], has been proven to be effective in tasks such as image classification, computer vision, and object detection [31]. The feature matrix can be likened to an image with its matrix structure, allowing it to be used as input to a CNN. A simple CNN was implemented for the experiments in this study, with two convolutional layers to downsample, two convolutional layers to upsample, and the ReLU activation function between each layer.

Feedforward Neural Network (FNN): The architecture of this model is composed of fully-connected layers as opposed to the convolutional layers of a CNN [32]. The FNN has found use in several machine learning applications due to its simplicity and computational efficiency. In FNNs, data flows in one direction from input to output, hence the name ”feedforward.” Similar to the CNN implementation mentioned above, this study uses a two-layer FNN, once again using the ReLU activation function between the two layers.

K-Nearest Neighbors (KNN): This simple algorithm is built on the premise that objects in the same neighborhood are likely to be classified in the same way. The algorithm, as the name suggests, uses the classes of the K nearest neighbors of a given object to determine its class [33]. KNN is used in many different ways, including text categorization, pattern recognition, and image classification. We employ a KNN classifier with $k = 5$.

Logistic Regression: This classifier uses a linear combination of inputs and their corresponding weights and performs a non-linear transformation on it using the sigmoid function. The model is used primarily for binary classification tasks due to the logistic data distribution of the sigmoid function [34]. The classifier is designed to determine decision boundaries among classes.

The baseline models cannot use the set notation of a graph directly as input since it does not consist of consistent features (the node embeddings of GNNs are instead calculated during GNN training). However, the graph $\mathcal{G}$ can be converted into a feature matrix by extracting features from the graph [35]. Let $\Gamma(u)$ represent the set of neighbors of node u and $k(u)$ represent the degree of node u. The following features are extracted from the topology of the graph to form the feature matrix for training our baseline models:

1. *Jaccard Coefficient* [36]: This coefficient is used for information retrieval, using both the intersection of common features and the union of all features in the graph. For graph $\mathcal{G}$, the features are the students with whom a given student u has interacted.

2. *Adamic-Adar (AA) Index* [36] [37]: This index computes node similarity by shared features, or neighbors, among nodes. Through use of the logarithmic function, nodes with fewer neighbors are given more weight than nodes with more neighbors.

3. *Resource Allocation (RA) Index* [36] [37]: In complex networks, resource allocation is promoted. For example, given a graph with two unconnected nodes u and v, the common neighbor of node u transmits, or allocates, a fraction of its resources to node v. This index minimizes the effect of nodes of higher degrees.

4. *Preferential Attachment (PA) Score* [36] [37]: This score measures the growth of the neighborhoods as the graph nodes are joined. As the name implies, nodes with higher degrees are more likely to join to other nodes than nodes with lower degrees.

5. *Cosine Similarity (CS) Index* [36] [37]: This metric, also referred to as the Salton Index, computes the cosine angle between the vectors of two nodes. A higher value of the Salton Index implies a stronger alignment between the neighborhoods of the two nodes.

6. *Dice Coefficient* [38]: This coefficient, similar to the Jaccard coefficient, uses the magnitude of the intersection of the neighbors of two given nodes. However, the Dice coefficient includes the number of overlapping neighbors in its denominator and also doubles the numerator, giving more weight to the common neighbors of the two nodes.

Each of the six aforementioned graph features is computed for each node pair, or edge, and placed in a feature vector. The feature matrix consists of the feature vectors for each node pair. Thus, for M edges, an $M \times 6$ feature matrix is formed.

Table 1: Performance metrics for each link prediction model. Here, we see that our proposed GNN framework performed the best with respect to AUC, outperforming each of the considered baselines. While the CNN model performed the best with respect to ACC, ACC is a biased metric for our imbalanced class case, thus corroborating our proposed GNN's best overall performance.

Training Set Through	Metric	GNN	GNN (bin)	CNN	FNN	KNN	LR
Week 1	AUC	0.5000	0.5000	0.5000	0.5000	0.5000	**0.5404**
	ACC	0.5007	0.3938	0.8731	0.8731	0.8731	0.2508
Week 2	AUC	0.4984	0.5000	0.5000	0.5000	**0.5362**	0.5037
	ACC	0.5000	0.3924	0.8641	0.8641	0.6667	0.1424
Week 3	AUC	0.5000	0.5000	0.5000	0.5000	0.5000	**0.5085**
	ACC	0.5036	0.3884	0.8571	0.8571	0.8571	0.7253
Week 4	AUC	0.5000	0.5000	0.5000	0.5000	0.4930	**0.5128**
	ACC	0.4907	0.3869	0.8664	0.8664	0.8543	0.8664
Week 5	AUC	**0.5462**	0.5000	0.5000	0.5000	0.5017	0.5152
	ACC	0.5737	0.3837	0.8692	0.1308	0.8458	0.8692
Week 6	AUC	**0.6625**	0.6108	0.5000	0.5000	0.4936	0.5000
	ACC	0.8297	0.3812	0.8814	0.8814	0.8701	0.8814
Week 7	AUC	**0.7479**	0.6273	0.5000	0.5000	0.5000	0.5000
	ACC	0.8284	0.6038	0.8913	0.8406	0.8913	0.8913
Week 8	AUC	**0.8000**	0.6223	0.5000	0.4689	0.5000	0.5000
	ACC	0.7774	0.7194	0.9293	0.9293	0.9293	0.9293

4.4 Performance Analysis

Here, the results of the experiments are reported and explained in terms of standard accuracy. The data used was highly biased toward scores of 9 and 10, and the effects of this skew will be seen in the performance analysis of the models.

Table 1 outlines our link prediction results, in terms of both considered metrics, in comparison to all baselines. Each model showed a general upward trend of accuracy over time. Although the GNNs achieved lower accuracy than all baseline models, producing final accuracies of 0.7774 and 0.7194 for the models with raw data and binarized data, respectively, while other models achieved final accuracies of 0.9293, accuracy was found to be an ineffective metric for measuring the performance of link prediction models due to the strong class inbalance of the input data (i.e., there were significantly more unformed links than formed links). Thus, we instead focus on AUC, which can better evaluate the performance of models in distinguishing between classes. Observing the progression of the AUC values over time, the baseline models all have AUC values around 0.5, which, as discussed before, is the same value as if the classifiers were forming random predictions. These AUC values do not improve over time but remain within 0.04 of the AUC value obtained from random selection. However, both GNNs have ROC curves that improve over time as more links are formed. The GNN trained on the raw data performed especially well, with a final AUC of 0.8000 by week 8, and the GNN trained on the binarized data performed better than the baseline models, with a final AUC of 0.6223. Notably, the raw data GNN was also the best-performing model after only four weeks of training data, demonstrating its efficacy even with little data to train on.

This study was conducted in a university that uses the quarter system, but based on the observation that the GNN performed better over time, the GNN may perform even better in a university that uses the semester system, where there is more time for the GNN to train on the evolving data. The GNN also demonstrated its effectiveness in a smaller classroom, even though this setting has less data. As mentioned above, the GNN outperformed all baseline models with the same data. This finding highlights the usefulness of the method proposed in this study for courses with lower enrollment or in smaller institutions where smaller classes are commonplace.

5 Conclusion

In this work, we developed a link prediction framework for social learning networks (SLNs) in in-person classroom settings by quantifying behavioral data via interaction surveys. Using this data, we not only showed a positive correlation between student learning interactions and academic outcomes, but also trained our link prediction framework, which outperformed multiple common baselines and demonstrated improved performance over time. Our framework can be used in classrooms to better understand the relationship between behavioral characteristics and course engagement, which has a known correlation with academic performance. Moreover, our framework can reveal engagement trends throughout the duration of a course, assisting instructors with course pacing as well as early and timely intervention for (i) individual students, (ii) groups of students, and (iii) potentially the entire class. Accurate link prediction also results in more beneficial collaborative learning experiences (i.e., students are methodically and efficiently paired with other students whose behavior and learning style are compatible with theirs), which is known to promote better learning outcomes. Furthermore, with the vast research on the importance of behavior over academic performance in student learning, our work provides a stronger indicator of academic success than previous academic measures, as our insights are drawn solely from behavioral data. Further studies will explore the use of this framework in different classroom settings, in terms of both classroom size and topic, to understand how this framework will perform in various courses. This study could also be extended to include additional behavioral data as well as implementations in K-12 settings.

Acknowledgments

This work was supported in part by the UC San Diego Academic Senate under grant RG116457 and in part by the National Science Foundation under grant ECCS-2512912.

References

[1] E. Alhazmi and A. Sheneamer, "Early predicting of students performance in higher education," *IEEE Access*, vol. 11, pp. 27 579–27 589, 2023.

[2] G. Van Helden, B. T. C. Zandbergen, M. M. Specht, and E. K. A. Gill, "Collaborative learning in engineering design education: A systematic literature review," *IEEE Transactions on Education*, vol. 66, no. 5, pp. 509–521, 2023.

[3] C.-C. Kiu, “Data mining analysis on student’s academic performance through exploration of student’s background and social activities,” in *2018 Fourth International Conference on Advances in Computing, Communication & Automation (ICACCA)*, 2018, pp. 1–5.

[4] L. Zhao, K. Chen, J. Song, X. Zhu, J. Sun, B. Caulfield, and B. M. Namee, “Academic performance prediction based on multisource, multifeature behavioral data,” *IEEE Access*, vol. 9, pp. 5453–5465, 2021.

[5] N. M. Webb, “Student interaction and learning in small groups,” *Review of Educational research*, vol. 52, no. 3, pp. 421–445, 1982.

[6] A. O. Ahmed Ismail and A. K. Mahmood, “The influence of interaction on students academic performance: Case of saudi public university,” in *2018 4th International Conference on Computer and Information Sciences (ICCOINS)*, 2018, pp. 1–6.

[7] Y. Zhen, J.-D. Luo, and H. Chen, “Prediction of academic performance of students in online live classroom interactions—an analysis using natural language processing and deep learning methods,” *Journal of Social Computing*, vol. 4, no. 1, pp. 12–29, 2023.

[8] S. R. Jayasekaran, S. Anwar, K. Cho, and S. F. Ali, “Relationship of students’ engagement with learning management system and their performance-an undergraduate programming course perspective,” in *2022 ASEE Annual Conference & Exposition*, 2022.

[9] F. H. Wang, “An exploration of online behaviour engagement and achievement in flipped classroom supported by learning management system,” *Computers & Education*, vol. 114, pp. 79–91, 2017.

[10] J. L. Harvey and S. A. P. Kumar, “A practical model for educators to predict student performance in k-12 education using machine learning,” in *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2019, pp. 3004–3011.

[11] B. A. Alnasyan, M. Basheri, and M. O. Alassafi, “A comprehensive comparative analysis of deep learning models for student performance prediction in virtual learning environments: Leveraging the oula dataset and advanced resampling techniques,” *IEEE Access*, vol. 13, pp. 75 953–75 972, 2025.

[12] M. Arashpour, E. M. Golafshani, R. Parthiban, J. Lamborn, A. Kashani, H. Li, and P. Farzanehfar, “Predicting individual learning performance using machine-learning hybridized with the teaching-learning-based optimization,” *Computer Applications in Engineering Education*, vol. 31, no. 1, pp. 83–99, 2023.

[13] V. Balachandar and K. Venkatesh, “Uncovering social dynamics in education through graph neural networks and continual learning methods,” in *2024 International Conference on Innovative Computing, Intelligent Communication and Smart Electrical Systems (ICSES)*, 2024, pp. 1–7.

[14] M. Laal and S. M. Ghodsi, “Benefits of collaborative learning,” *Procedia - Social and Behavioral Sciences*, vol. 31, pp. 486–490, 2012, world Conference on Learning, Teaching & Administration - 2011. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877042811030205

[15] T. Panitz and P. Panitz, *Encouraging the Use of Collaborative Learning in Higher Education*, 10 2018, pp. 161–202.

[16] T. Nokes-Malach, J. Richey, and S. Gadgil, "When is it better to learn together? insights from research on collaborative learning," *Educational Psychology Review*, vol. 27, 06 2015.

[17] S. Anwar, A. A. Butt, and M. Menekse, "Exploring relationships between academic engagement, application engagement, and academic performance in a first-year engineering course," in *2022 IEEE Frontiers in Education Conference (FIE)*, 2022, pp. 1–5.

[18] Y. Li, L. Qiu, and B. Sun, "School engagement as a mediator in students' social relationships and academic performance: a survey based on citespace," *International Journal of Crowd Science*, vol. 5, no. 1, pp. 17–30, 2021.

[19] J. M. Fletcher, S. L. Ross, and Y. Zhang, "The consequences of friendships: Evidence on the effect of social relationships in school on academic achievement," *Journal of Urban Economics*, vol. 116, p. 103241, 2020.

[20] W. Zou, X. Hu, Z. Pan, C. Li, Y. Cai, and M. Liu, "Exploring the relationship between social presence and learners' prestige in mooc discussion forums using automated content analysis and social network analysis," *Computers in Human Behavior*, vol. 115, p. 106582, 2021.

[21] D. Das, "A novel negative link prediction algorithm for social networks," in *2022 IEEE 95th Vehicular Technology Conference: (VTC2022-Spring)*, 2022, pp. 1–5.

[22] Y. Liu, Y. Jiang, P. Shen, and J. Jia, "Link prediction in directed graphs via dummy node line graphs," in *2024 5th International Conference on Artificial Intelligence and Computer Engineering (ICAICE)*, 2024, pp. 533–537.

[23] X. Liu, J. Cheng, Y. Song, and X. Jiang, "Boosting graph structure learning with dummy nodes," in *International conference on machine learning*. PMLR, 2022, pp. 13 704–13 716.

[24] Z. Zhang, X. Wu, G. Zhu, W. Qin, and N. Liang, "A graph attention network-based link prediction method using link value estimation," *IEEE Access*, vol. 12, pp. 34–45, 2024.

[25] T. Rubasinghe, S. Vasanthapriyan, and G. Romiyal, "Ai-driven framework for enhancing student engagement through real-time monitoring of attendance and cognitive focus in physical classrooms," in *2025 International Research Conference on Smart Computing and Systems Engineering (SCSE)*, 2025, pp. 1–8.

[26] W. L. Hamilton, R. Ying, and J. Leskovec, "Inductive representation learning on large graphs," 2018. [Online]. Available: https://arxiv.org/abs/1706.02216

[27] M. Abedini and H. Shakibian, "Inter-layer similarity-based graph neural network for link prediction in social multiplex networks," in *2024 10th International Conference on Web Research (ICWR)*, 2024, pp. 76–80.

[28] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017. [Online]. Available: https://arxiv.org/abs/1412.6980

[29] D. Meghana and M. A, “Rcs-yolo: A rapid and precise brain tumor detection framework utilizing the adam optimizer,” in *2025 International Conference on Ambient Intelligence in Health Care (ICAIHC)*, 2025, pp. 1–6.

[30] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.

[31] K. Selvarajah, K. Ragunathan, Z. Kobti, and M. Kargar, “Dynamic network link prediction by learning effective subgraphs using cnn-lstm,” in *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020, pp. 1–8.

[32] N. Zhang, W. Wu, and G. Zheng, “Convergence of gradient method with momentum for two-layer feedforward neural networks,” *IEEE Transactions on Neural Networks*, vol. 17, no. 2, pp. 522–525, 2006.

[33] T. Mladenova and I. Valova, “Comparative analysis between the traditional k-nearest neighbor and modifications with weight-calculation,” in *2022 International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT)*, 2022, pp. 961–965.

[34] N. R, B. R, R. Balamanigandan, and A. Priscilla, “Analyzing the performance of novel logistic regression over linear regression algorithms for predicting fake job with improved accuracy,” in *2024 5th International Conference on Electronics and Sustainable Communication Systems (ICESC)*, 2024, pp. 1728–1732.

[35] A. P. Hridi, M. Hoq, Z. Gao, C. Lynch, R. Sahay, S. Hosseinalipour, and B. Akram, “Privacy-preserving distributed link predictions among peers in online classrooms using federated learning,” 2025. [Online]. Available: https://arxiv.org/abs/2504.10456

[36] A. JIBOUNI, D. LOTFI, M. E. MARRAKI, and A. HAMMOUCH, “A novel parameter free approach for link prediction,” in *2018 6th International Conference on Wireless Networks and Mobile Communications (WINCOM)*, 2018, pp. 1–6.

[37] S. U. Balvir, M. M. Raghuwanshi, and P. D. Shobhane, “An overview of similarity-based methods in predicting social network links: A comparative analysis,” *IEEE Access*, vol. 12, pp. 120 913–120 934, 2024.

[38] S. Abinash, “Two novel graph theory based algorithms for salient object detection,” in *2018 4th International Conference on Computing Communication and Automation (ICCCA)*, 2018, pp. 1–6.