

Spatiotemporal Link Formation Prediction in Social Learning Networks Using Graph Neural Networks

Ali Mohammadiasl
UC San Diego
amohammadiasl@ucsd.edu

Seyyedali Hosseinalipour
University at Buffalo–SUNY
alipour@buffalo.edu

Bitu Akram
NC State University
bakram@ncsu.edu

Rajeev Sahay
UC San Diego
r2sahay@ucsd.edu

ABSTRACT

Social learning networks (SLNs) are graphical representations that capture student interactions within educational settings (e.g., a classroom), with nodes representing students and edges denoting interactions. Accurately predicting future interactions in these networks (i.e., link prediction) is crucial for enabling effective collaborative learning, supporting timely instructional interventions, and informing the design of effective group-based learning activities. However, traditional link prediction approaches are typically tuned to general online social networks (OSNs), often overlooking the complex, non-Euclidean, and dynamically evolving structure of SLNs, thus limiting their effectiveness in educational settings. In this work, we propose a graph neural network (GNN) framework that jointly considers the temporal evolution within classrooms and spatial aggregation across classrooms to perform link prediction in SLNs. Specifically, we analyze link prediction performance of GNNs over the SLNs of four distinct classrooms across their (i) temporal evolutions (varying time instances), (ii) spatial aggregations (joint SLN analysis), and (iii) varying spatial aggregations at varying temporal evolutions throughout the course. Our results indicate statistically significant performance improvements in the prediction of future links as the courses progress temporally. Aggregating SLNs from multiple classrooms generally enhances model performance as well, especially in sparser datasets. Moreover, we find that jointly leveraging both the temporal evolution and spatial aggregation of SLNs significantly outperforms conventional baseline approaches that analyze classrooms in isolation. Our findings demonstrate the efficacy of educationally meaningful link predictions, with direct implications for early-course decision-making and scalable learning analytics in and across classroom settings.

Keywords

Graph Neural Networks, Link Prediction, Social Learning Networks, Social Network Analysis.

Ali Mohammadiasl, Bitu Akram, Seyyedali Hosseinalipour, and Rajeev Sahay. Spatiotemporal Link Formation Prediction in Social Learning Networks Using Graph Neural Networks. In Anthony Botelho, Maria Mercedes T. Rodrigo, Adish Singla, Hiroaki Ogata, Hyejeong So, and Young Hoan Cho (eds.) Proceedings of the 19th International Conference on Educational Data Mining, Seoul, Republic of Korea, June, 2026, pp. 541–548. International Educational Data Mining Society (2026).

© 2026 Copyright is held by the author(s). This work is distributed under the Creative Commons Attribution NonCommercial NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.

<https://doi.org/10.5281/zenodo.21040072>

1. INTRODUCTION

Collaborative learning is a major pillar of modern education, shaping student engagement, persistence, and academic success. Consequently, instructors increasingly incorporate structured peer interactions, such as peer instruction and discussion forums, to enhance learning outcomes [1, 3]. Yet, in practice, collaborative behaviors do not emerge uniformly across students or over time: some students rapidly become socially central, while others remain weakly connected or isolated, particularly in the early stages of a course. These uneven and evolving interaction patterns complicate instructors' ability to anticipate collaboration dynamics and intervene effectively when support is most needed.

In order to model and analyze such interaction dynamics, social learning networks (SLNs) have emerged as a powerful abstraction. SLNs represent classroom interactions as graphs, where *nodes* correspond to students and *edges* denote interactions arising from discussions, question–answer exchanges, or collaborative activities [12, 22]. This abstraction enables a principled investigation of network inference problems in educational graphs, with *link prediction* (i.e., forecasting which student interactions are likely to form in the future) serving as one of the fundamental problems in this setting.

1.1 Challenges and Motivations

Within SLNs, *link prediction* has emerged as a valuable tool for forecasting future collaborations, identifying at-risk students, and informing instructional interventions. Furthermore, accurate link prediction can reveal relationships between students' behavioral characteristics and their tendencies toward class involvement, which is often linked to their learning performance [29]. However, prior work on link prediction in educational settings has largely focused on analyzing each classroom independently, treating the SLN associated with a single course offering as an isolated graph [20, 30]. While this approach is intuitive, it implicitly assumes that meaningful structure must be learned anew for every classroom, ignoring common patterns of social and pedagogical dynamics that recur across courses, instructors, and academic terms. In addition, this approach fails to capture the temporal processes through which collaborations form early in the course when few observed links have been formed (i.e., the cold-start problem), thus preventing accurate link prediction precisely when early intervention would be most beneficial. Moreover, the inherent complexity of SLNs, including their evolving topology and irregular non-Euclidean

structure, renders traditional data-driven models [19] [27], such as convolutional neural networks (CNNs), ill-equipped for the task of link prediction.

Graph Neural Networks (GNNs), on the other hand, offer a powerful alternative to traditional CNNs when it comes to predicting links in SLNs. Unlike CNNs, which are primarily designed for grid-like data such as images, GNNs are tailored for irregular graph structures. This makes them particularly well-suited for SLNs, where nodes (students) and edges (interactions) form complex non-Euclidean structures [5]. Despite their potential, the application of GNNs in SLNs remains in its early stages, with only a handful of studies exploring their potential and limitations [20].

To address this gap in the literature, *our work leverages GNNs to unlock two critical domains in link prediction for SLNs: the ‘temporal’ domain and the ‘spatial’ domain.* Specifically, in the temporal domain, we examine the evolution of SLN graphs over time, while in the spatial domain, we investigate the effects of combining SLN graphs from different classroom contexts to improve link prediction performance equitably across classrooms.

1.1.1 Exploring the Temporal Domain

SLNs are inherently dynamic: the topology of these networks evolves as student interactions change throughout the course. This temporal evolution can significantly affect GNNs’ future link prediction performance in SLNs. By analyzing how SLN structures shift at different stages of a course, we can uncover deeper insights into student interactions and the conditions under which link prediction performs best. We subsequently present and investigate, for the first time in the literature, the following hypothesis:

Hypothesis 1: *Do SLNs’ temporal structures have a significant impact on the prediction capability of GNNs trained on them? If yes/no, how do link prediction performance metrics vary during the temporal evolution of SLNs?*

1.1.2 Investigating the Spatial Domain

Our study on the spatial domain of SLNs introduces the combination of SLN graphs across different classrooms. At first glance, merging SLNs might seem to improve model performance by providing more data. However, our intuition suggests that the unique topological structures inherent in each SLN, stemming from varied interaction patterns among peers in different classrooms, can lead to unexpected outcomes. Subsequently, our formally proposed hypothesis that we investigate for the first time in the literature is as follows:

Hypothesis 2: *Does combining SLNs together lead to better performance of the GNN models trained on them? If yes/no, how do link prediction performance metrics vary upon combining SLNs?*

1.1.3 Spatiotemporal Analysis

A careful observation of the interrelations between the temporal and spatial domains reveals non-trivial interactions. For instance, our intuition may suggest that combining SLNs at earlier stages of a course might be more beneficial, as networks are still forming and less defined. However, as SLNs develop their distinct patterns later in the course, merging them might not yield the same performance improvements.

To encapsulate this nuanced view, we propose and investigate the following hypothesis:

Hypothesis 3: *Does combining SLNs together at different time stamps during a course lead to varying performance of the GNN models trained on them? If yes/no, how do link prediction performance metrics vary during the combination of SLNs at different time stamps of a course?*

1.2 Overview and Summary of Contributions

Formally, our contributions can be summarized as follows:

- We pioneer an investigation into the temporal evolution and spatial aggregation of SLNs by analyzing how link formation and interaction patterns change throughout a class and how classrooms can benefit one another through sharing their data. Our time-evolving analysis over multiple classroom SLNs provides novel insights into these dynamics.
- We develop a GNN-based framework designed to tackle the unique challenges presented by sparse, imbalanced, and temporally evolving SLN datasets. This framework is optimized to capture the complex topological features inherent in SLNs. We then use this framework to validate our ***Hypothesis 1.***
- We demonstrate that combining data from multiple SLNs, spanning different classroom contexts, can enhance link prediction performance under certain conditions. This contribution is focused on validating our ***Hypothesis 2.***
- By examining the interplay between the temporal evolution and spatial data aggregation of SLNs, we uncover non-trivial interactions that influence overall link prediction accuracy. This integrated analysis addresses our ***Hypothesis 3,*** where our findings offer a refined understanding of how temporal and spatial factors collectively inform link prediction strategies.

2. RELATED WORK

Link prediction in general computational graphs (i.e., non-SLN graphs such as Online Social Networks (OSNs)) has evolved through various methodologies over time. Early methods relied on deriving features from the graph topology [15, 17] to infer future trends [13, 31]. While effective in some cases, these heuristic-based methods often fail to capture the inherent complexities of dynamic, rapidly-evolving graphs. To address these limitations, more rigorous link prediction approaches were developed, such as matrix factorization and probabilistic models, which aim to learn low-dimensional representations of nodes [11], [26], [18], [33]. Although these models yielded improved prediction accuracy, later methods further advanced performance by offering greater scalability and adaptability to capture the complex, evolving nature of OSNs. In particular, the advent of feature-based deep learning introduced models based on artificial neural networks for link prediction, which leverage node characteristics and graph adjacency information to improve link prediction in both general OSNs [8, 16, 24, 32] as well as SLNs [7, 20, 30]. While these methods achieved notable success, they typically treat graphs as static structures and lack the capability to effectively capture the temporal dynamics inherent in SLNs. Moreover, they fail to account for

the intricate topological structures and co-temporal evolution of interactions across multiple classrooms, limiting their applicability in educational settings.

More recently, GNNs have emerged as powerful tools for link prediction in general graph structures due to their ability to learn both the local structure and the global topology of a graph [2, 4, 9, 14, 21, 23, 28]. GNN-based models, such as Graph Convolutional Networks (GCNs) [10], GraphSAGE [6], and Graph Attention Networks (GATs) [25], have demonstrated superior performance compared to feature-based deep learning models by aggregating neighborhood information through learned message-passing schemes among nodes. Despite these advancements, no prior work has specifically examined GNN-based link prediction approaches in dynamically evolving SLNs, thus overlooking the temporal evolution and spatial aggregation properties that are unique to SLNs. As a result, the impact of SLN dynamics over time, the integration of information across classroom contexts, and the interplay between these two dimensions remain largely unexplored. Our work aims to address this critical gap in the literature.

3. METHODOLOGY

In this section, we begin by describing the mathematical representation of SLNs as graphs (Sec. 3.1). Next, we outline the task of link prediction in SLNs, detailing how GNNs are used to infer future links (Sec. 3.2). We then introduce our joint classroom approach, detailing the integration of SLNs to improve predictive performance in individual classrooms (Sec. 3.3). Lastly, we describe the evaluation metrics used to assess the link prediction performance (Sec. 3.4).

3.1 Representation of SLNs as Graphs

We represent each SLN as a graph $G = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} denotes the set of *nodes* (i.e., students) and \mathcal{E} represents the set of *edges* (i.e., interactions between students). Each edge $e_{ij} \in \mathcal{E}$ signifies an interaction between two nodes/students v_i and v_j , where $v_i, v_j \in \mathcal{V}$, such as collaborative projects, discussions, or information exchanges. Subsequently, we define the adjacency matrix of the SLN graph as $\mathbf{A} = [A_{ij}]_{v_i, v_j \in \mathcal{V}}$, where

$$A_{ij} = \begin{cases} 1, & \text{if } (v_i, v_j) \in \mathcal{E} \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

In words, A_{ij} represents the presence ($A_{ij} = 1$) or absence ($A_{ij} = 0$) of a connection/edge between nodes v_i and v_j .

3.2 Link Prediction in SLNs

The temporal evolution of an SLN is illustrated in Fig. 1, showcasing a simplified example of an SLN across three sequential stages of a course. Subsequently, given the SLN graph described in Sec. 3.1 at some point t during the duration of a course/class (e.g., week 4 of the semester), our objective is to predict which links will form at a future time $t+t_0$ (e.g., week 6) between students by using GNNs. To this end, we adopt GraphSAGE [6], which learns node embeddings/features by iteratively aggregating information from local neighborhoods. As new nodes are not introduced in the SLN, this inductive learning approach enables the model to dynamically adapt node embeddings as local structures (i.e., student interaction patterns) evolve over time, making it well-suited for the temporal changes of student interactions throughout the course.




Figure 1: Visualization of SLN temporal evolution. Left to right shows the progress of the classroom from the beginning to the end. A solid line denotes an existing edge at that stage, and a dotted line denotes a non-existing edge at that stage that will become an existing one in a future stage.




Figure 2: An overview of the GraphSAGE pipeline for link prediction. Left: Sampling stage, where the 1-hop and 2-hop neighborhoods of the target nodes (highlighted in green) are collected. Middle: Aggregation stage, where each node’s embedding is updated by aggregating the features of its neighbors using multiple layers of aggregation (in our case, mean aggregator). Right: Prediction stage, where a dot product is computed between the final embeddings of the two target nodes to predict the existence of a link.

3.2.1 Link Predictions based on GNNs in a Nutshell

In a nutshell, to predict whether an interaction (i.e., a link) will form between two students v_i and v_j , we compute a numerical representation (embedding) for each student using a GraphSAGE model as we will later discuss (see Sec. 3.2.2). These embeddings capture the structural context of each student within the SLN. The likelihood of a future interaction is then estimated by taking the dot product of the respective node embeddings (potentially after propagating through one or more hidden layer) as follows:

$$\hat{y}_{v_i, v_j} = \sigma(h_{v_i}^\top h_{v_j}), \quad (2)$$

where $\sigma(\cdot)$ is the sigmoid function, h_{v_i} and h_{v_j} are the learned embeddings for students v_i and v_j , respectively ($h_{v_i}, h_{v_j} \in \mathbb{R}^d$), and $\hat{y}_{v_i, v_j} \in \mathbb{R}$.

3.2.2 Obtaining Node Embeddings via GNN

To obtain the node embeddings, we use a *two-layer GraphSAGE model*, where each layer aggregates information from a student’s neighbors in the SLN. This model, depicted in Fig. 2, replicates a convolutional filter that is capable of operating in a non-Euclidean space (i.e., graph topology). Intuitively, at each SAGEConv layer, GraphSAGE updates the representation of a node by combining its own features with the aggregated features of its neighbors. By the second layer, the embedding of a node reflects not only its own character-




Figure 3: Visualization of spatial evolution. The left displays a single SLN A (representing a single classroom), the middle is SLN A combined with SLN B (i.e., two classrooms combined together), and the right is SLN A and B combined with SLN C (i.e., three classrooms combined together).

istics but also the structure and features of its broader local neighborhood, enabling the model to better infer potential future connections among nodes. In particular, the model begins with an input feature matrix $\mathbf{X} = [x_{v_i}]_{v_i \in \mathcal{V}} \in \mathbb{R}^{|\mathcal{V}| \times F}$, where the i^{th} row (i.e., $x_{v_i} \in \mathbb{R}^F$) corresponds to the feature vector of student v_i consisting of F features. Since our SLNs do not contain explicit node features (e.g., profile or grade data), each student is initialized with a unique trainable embedding vector. These initial embeddings are then refined through neighbor aggregation as follows.¹

In the first layer, the model updates each student’s embedding by averaging its own feature vector with those of its directly connected peers (1-hop neighbors). This step captures local interaction patterns through the following embedding:

$$h_{v_i}^{(1)} = \zeta \left(W_1 \cdot \text{MEAN} \left(\{x_{v_j} | v_j \in \mathcal{N}(v_i) \cup \{v_i\}\} \right) + b_1 \right), \quad (3)$$

where $\zeta(\cdot)$ is the activation function, $h_{v_i}^{(1)} \in \mathbb{R}^{d_1}$ is the updated embedding for student v_i after one round of aggregation, $\mathcal{N}(v_i)$ is the set of students directly connected to v_i in the SLN (i.e., students who have interacted with v_i), $x_{v_j} \in \mathbb{R}^F$ is the feature vector of student v_j , W_1 and b_1 are trainable weights and bias parameters, and $\text{MEAN}(\cdot)$ takes the element-wise average across the set.

In the second layer, the model does the same aggregation, except this time, it uses the embedding from the first layer, allowing to implicitly capture information from 2-hop neighbors. In particular, the second layer of the model captures the following embedding:

$$h_{v_i}^{(2)} = \zeta \left(W_2 \cdot \text{MEAN} \left(\{h_{v_j}^{(1)} | v_j \in \mathcal{N}(v_i) \cup \{v_i\}\} \right) + b_2 \right), \quad (4)$$

where $h_{v_i}^{(2)}$ is the final embedding for student v_i (i.e., used as h_{v_i} in (2)), and W_2 and b_2 are trainable parameters.

3.3 Classroom Merging and GNN Training

3.3.1 Merging Classrooms

To assess the impact of training a GNN for link prediction, we conduct experiments on individual and joint SLNs, capturing

¹Fig. 2 illustrates GraphSAGE assuming aggregation over all neighbors for clarity. However, in practice, GraphSAGE samples a fixed-size subset of neighbors at each layer to reduce computational overhead and memory usage.

their spatial evolution. In particular, initially, we train separate GNN models for each classroom, using the same time instances of the SLN graph across the classrooms to ensure consistency across experiments. We then construct combined SLN graphs by merging classroom SLNs, systematically evaluating the combination of multiple classroom SLNs into a single joint graph. Subsequently, we train GNNs on these combined SLNs and evaluate their performance on individual classrooms to verify the impact of classroom combination on the link prediction performance of the trained models. We visualize the spatial aggregation of SLNs in Fig. 3.

3.3.2 GNN Training

To train GNNs over our SLNs (either on isolated or merged classrooms), we first note that there is an extreme class imbalance inherent in SLNs. This imbalance arises because *SLNs are sparse graphs*: the vast majority of possible node pairs in the graph do not share an edge, meaning that most students do not interact with each other.

To mitigate this imbalance during training, we use a weighted binary cross-entropy loss: for a predicted link score s of a given node pair and the respective ground truth label $y \in \{0, 1\}$ (1 for edge, 0 for non-edge), the loss is defined as:

$$\ell(s, y) = \begin{cases} \text{PCW} \cdot \log(1 + e^{-s^+}), & \text{if } y = 1, \\ \log(1 + e^{s^-}), & \text{if } y = 0. \end{cases}$$

The total loss across a positive and negative sample pair is:

$$\mathcal{L} = \text{PCW} \cdot \log(1 + e^{-s^+}) + \log(1 + e^{s^-}),$$

where s^+ and s^- denote the scores for a positive (i.e., connected) and negative (i.e., not connected) student pair, respectively. We found through a grid search, optimizing for the minimization of false negatives, that fixing *positive class weight* $\text{PCW} = 20$ provides a stable and effective training regime across all SLNs, balancing the impact of class imbalance without over-amplifying the minority class loss.

3.4 Evaluation Metrics

Although accuracy is a conventional metric for classification tasks, it is not appropriate in link prediction due to the large class-imbalance, where a model with relatively high accuracy may fail to predict any edges. Instead, we use the *Area Under the receiver operating characteristic Curve (AUC)*, as it provides a robust measure of predictive effectiveness in class-imbalanced settings since it evaluates a model’s ability to rank positive samples (i.e., student pairs with edges) higher than negative ones, making it a more reliable indicator of performance in sparse graphs. The AUC is computed by

$$\text{AUC} = \frac{1}{|\mathcal{P}||\mathcal{N}|} \sum_{(s^+, s^-) \in (\mathcal{P}, \mathcal{N})} \mathbb{I}(s^+ > s^-), \quad (5)$$

where \mathcal{P} and \mathcal{N} denote the sets of predicted scores for positive and negative samples, respectively, and $\mathbb{I}(\cdot)$ is an indicator function that returns 1 if the positive score s^+ is greater than the negative score s^- , and 0 otherwise. In our context, the model assigns a score to each link (i.e., a student pair), where a higher score indicates a higher likelihood that the link should exist and a threshold score of 0.5 is used to determine the prediction of a link.

Table 1: Metrics for each considered dataset including the number of nodes, edges, duration (weeks), and posts.

Dataset	# Nodes	# Edges	Duration	# Posts
VS	677	4702	5	7484
ML	3290	30610	12	25481
AL	1165	6773	13	16276
CP	900	3418	8	8255

4. RESULTS AND DISCUSSION

Henceforth, we discuss our empirical setup (Sec. 4.1) followed by analyzing our results in temporal (Sec. 4.2), spatial (Sec. 4.3), and spatiotemporal settings (Sec. 4.4). The code used in this work is available at <https://github.com/Alimohammadiasl/GNNs-for-SLNs>.

4.1 Empirical Setup

For our empirical analysis, we employ four SLNs collected from four distinct Massive Open Online Course (MOOC) datasets: Virtual Shakespeare (VS), Machine Learning (ML), Algorithms (AL), and English Composition (CP). The SLN data we used was introduced in [30] and captures structured interactions that occurred between students on discussion forums. Specifically, if student v_i and v_j interact on a discussion forum by posting an answer or a follow-up question to an initial post, we consider $A_{ij} = 1$. Our considered datasets intentionally span both STEM and humanities courses so that the efficacy of our approach can be evaluated across various disciplines. Tab. 1 summarizes these datasets.

Our GraphSAGE model consists of an $F = 16$ -dimensional embedding for each node at the input followed by two convolutional layers. We use Rectified Linear Unit (ReLU) as the activation function following both hidden layers. For training, we perform full-graph updates for a total of 1500 epochs, optimizing the model with the Adam optimizer at a learning rate of $5e - 4$ using the above-described weighted binary-cross entropy loss function.

To ensure robustness (in terms of the time that the model is used during the term/semester), we evaluate our models in multiple temporal snapshots of the SLNs. We define progress as the percentage of the term that has elapsed, with 0% representing the start and 100% the end. We consider 25%, 50%, 75%, and 90% progress, and we perform temporal k -fold cross-validation (with $k = 10$) and report the average AUC across all partitions for statistical validation.

4.2 Temporal Graph Analysis

We begin by examining how link prediction in SLNs evolve temporally throughout the duration of a single course (addressing **Hypothesis 1** from Sec. 1). Inspecting Tab. 2 and focusing on isolated GNN training (i.e., ‘GNN iso’), which represents GNNs trained on a particular course in isolation, we observe consistent AUC improvements as more temporal data become available in each course. As expected, the AUC improves steadily as the course progresses, and the model benefits from richer structural information in later stages, particularly for smaller courses (e.g., VS, AL, and CP), which suffer from link formation data early on, than larger ones

Table 2: AUC performance of each considered model across four stages of course progress (25%, 50%, 75%, 90%). Temporal trends can be observed across progress stages of each course, while spatial effects are seen by comparing ‘GNN iso’ and ‘GNN all’ at each progress stage.

Dataset	Progress	GNN all	GNN iso	GCN [10]	GAT [25]	CNN [20]
VS	25%	0.88	0.86	0.82	0.87	0.64
	50%	0.92	0.89	0.85	0.89	0.62
	75%	0.94	0.88	0.86	0.90	0.60
	90%	0.94	0.88	0.85	0.89	0.67
ML	25%	0.85	0.84	0.76	0.82	0.72
	50%	0.86	0.85	0.70	0.83	0.83
	75%	0.86	0.84	0.69	0.83	0.90
	90%	0.88	0.83	0.87	0.84	0.81
AL	25%	0.92	0.92	0.83	0.89	0.82
	50%	0.94	0.93	0.86	0.90	0.91
	75%	0.94	0.93	0.86	0.92	0.93
	90%	0.95	0.94	0.86	0.92	0.93
CP	25%	0.86	0.84	0.82	0.83	0.66
	50%	0.90	0.84	0.85	0.85	0.71
	75%	0.89	0.85	0.83	0.85	0.80
	90%	0.91	0.88	0.85	0.85	0.60

(ML). This indicates that the model becomes increasingly capable of distinguishing between likely and unlikely links between students as it observes more interactions. These trends indicate that temporal progress has a measurable impact on model performance, affirming **Hypothesis 1** from Sec. 1. Overall, later snapshots provide more robustness for learning, while early-course graphs pose prediction challenges due to the limited structural context.

4.3 Spatial Graph Analysis

We next analyze the spatial evolution of link prediction in SLNs by exploring how combining SLNs across different classrooms affects model performance (addressing **Hypothesis 2** from Sec. 1). To evaluate this, in Tab. 2, we compare the performance of the models when trained on individual SLNs versus on combined SLNs (i.e., ‘GNN iso’ and ‘GNN all’). The results reveal how merging SLNs positively impacts model performance across different stages of the course. For example, the AUC for VS improves from an average of 0.88 to an average of 0.94 when combined with others over the course progression. Moreover, our results show that, although combining SLNs generally improves AUC performance, particularly in smaller or sparser datasets like VS and CP, the gains are less pronounced in larger classrooms such as AL and ML, where the SLN topological structure is already rich and adding more data to it may not affect its topological significance in a meaningful way. This suggests that the benefits of spatial aggregation are not solely due to having more data, but rather depend on the compatibility of interaction patterns across classrooms. These results support **Hypothesis 2** from Sec. 1, affirming and revealing the importance and timing of spatial aggregation.

4.4 Spatiotemporal Graph Analysis

Table 3: Hypothesis testing (one-sided t -test, $\alpha = 0.10$) comparing GNN isolated (IS) vs. GNN combined (CO) vs. CNN baseline (CNN) at 25% and 75% course progress. Each row compares spatial performance, while reading down the table reveals temporal trends across progress stages.

Models	Datase	Progress	p -value	Decision
IS vs CO	AL	25%	$4.18e-1$	Fail to Reject H_0
IS vs CO	AL	75%	$8.03e-2$	Reject H_0
CNN vs CO	AL	25%	$6.02e-3$	Reject H_0
CNN vs CO	AL	75%	$1.25e-1$	Fail to Reject H_0
IS vs CO	CP	25%	$2.90e-2$	Reject H_0
IS vs CO	CP	75%	$1.42e-2$	Reject H_0
CNN vs CO	CP	25%	$5.75e-8$	Reject H_0
CNN vs CO	CP	75%	$1.23e-8$	Reject H_0
IS vs CO	VS	25%	$6.86e-4$	Reject H_0
IS vs CO	VS	75%	$8.00e-6$	Reject H_0
CNN vs CO	VS	25%	$8.13e-7$	Reject H_0
CNN vs CO	VS	75%	$3.7e-5$	Reject H_0
IS vs CO	ML	25%	$3.06e-1$	Fail to Reject H_0
IS vs CO	ML	75%	$6.14e-2$	Reject H_0
CNN vs CO	ML	25%	$5.77e-2$	Reject H_0
CNN vs CO	ML	75%	$2.15e-1$	Fail to Reject H_0

In Secs. 4.2 and 4.3, we saw that (i) SLNs become structurally richer as a course progresses (temporal evolution) and (ii) combining multiple classroom graphs (spatial evolution) generally boosts AUC by exposing the model to a broader distribution of link patterns. However, these two dimensions do not operate independently. In particular, the benefit of combining SLNs (i.e., using spatial features of the graph) appears to depend on when in the course (temporally) the graphs are merged. Thus, in this section, we address **Hypothesis 3** from Sec. 1, which investigates whether merging SLNs at different stages of course progress leads to measurable gains in model performance.

As shown in Tab. 2, we see that combining SLNs earlier in the course leads to higher performance improvements compared to merging them at later stages. This is consistent with the intuition that early in the term, when individual SLNs are sparse, combining information across classrooms helps stabilize learning and improves performance. In contrast, in the later stages of a course, each SLN develops distinct structural characteristics, potentially reflecting subject matter, classroom dynamics, or group behavior and, as a result, attain lower performance gains upon being combined with other SLNs towards the end of the course.

Further, to highlight the superiority of our methodology, we compare our GNN-based framework to the current state-of-the-art framework for link prediction in SLNs, which consists of manually engineering fixed graph features and using them as input into a convolutional neural network (CNN) for link prediction [20], along with state-of-the-art approaches for link prediction in general OSNs, which consist of a sum pooling graph convolutional network (GCN) [10] and a Graph Attention Network (GAT) [25]. As can be seen from Tab. 2, each baseline struggles to attain consistently high performance

across the spatiotemporal topology of the graph. The CNN in particular struggles due to its fixed feature engineering, whereas the GNN is able to offer high link prediction performance due to its dynamic adaptation to the SLN graph topology (via its node embeddings) and its evolution over time. This is highlighted in Tab. 2, where the GNN trained on a single SLN (GNN iso) generally outperforms the CNN baseline across all datasets and time instances. Furthermore, the GNN trained on combined SLNs (GNN all) not only outperforms the CNN but also significantly surpasses the isolated GNN as well as the other baselines in many cases, particularly in earlier stages of small courses. This demonstrates the dual benefit of our framework: first, to outperform traditional feature-based methods, and second, to exploit shared link formation patterns across SLNs to enhance predictive performance.

To further support the above observations, we perform formal hypothesis testing using a one-sided t test ($\alpha = 0.10$) comparing the AUC performance of our GNN model trained in combination (CO) with other courses against both the GNN trained in isolation (IS) on a single course and the CNN baseline (CNN). Here, we focus on the CNN baseline to demonstrate the statistical significance of using our framework’s learned features via node embeddings in comparison to fixed feature engineering as currently done in state of the art link prediction frameworks [20]. The null hypothesis in each test states that CO does not outperform the baseline, while the alternative asserts that CO achieves a higher AUC. Tab. 3 summarizes the results in 25% and 75% course progress for the four SLNs. We show results at these two progress points to highlight the superiority of our method in an early and a late stage of the course. Here, we see that the improvement of our framework is statistically significant, particularly for small courses (e.g., CP and VS). Larger courses (e.g., AL and ML), on the other hand, contain sufficient training data that can be captured in varying time instances. This is largely due to the bias of the larger courses as their SLNs account for a larger partition of the spatially aggregated graph structure. Yet, we still observe statistically significant gains in AUC when comparing CO against both CNN and IS. In particular, the p -values are especially low for smaller and sparser datasets such as VS and CP, even at early progress stages. This further supports **Hypothesis 3**: the spatiotemporal consideration of SLNs generally yields stronger link prediction performance compared to the current state of the art, particularly in smaller courses (CP and VS) that suffer from an abundance of links early in the course.

5. CONCLUSION

In this work, we developed a link prediction framework in SLNs through the lens of two critical dimensions: temporal evolution and spatial aggregation. Using multiple SLNs at varying progress stages, we analyzed how SLNs develop over time and how combining multiple SLNs affects model performance. Our results show that model performance improves as the SLNs progress through the course timeline: richer interaction structures at later stages offer stronger signals for link prediction. We also found that merging SLNs from multiple classrooms can enhance model performance as spatial aggregation helped overcome data sparsity in smaller courses with lower numbers of links throughout the course progression. Finally, our analysis of the spatiotemporal

aspects revealed that the benefits of merging graphs are most pronounced early in the course, and particularly in smaller courses, which require more structure early on.

6. ACKNOWLEDGMENT

This work was supported in part by the U.S. National Science Foundation (NSF) under Grant No. SaTC-2513164, ECCS-2512911, ECCS-2543754, and DRL-2418658.

7. REFERENCES

- [1] T. Babić, L. Kolar, and M. Miličević. Individual, cooperative and collaborative learning and students' perceptions of their impact on their own study performance. In *Proceedings of 44th International Convention on Information, Communication and Electronic Technology (MIPRO)*, pages 864–869, 2021.
- [2] X. Chen, P. I. Lei, Y. Sheng, Y. Liu, and Z. Gong. Social influence learning for recommendation systems. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 312–322, 2024.
- [3] F. P. A. de Medeiros and A. S. Gomes. An approach based on social network analysis to enhance social presence in a collaborative learning environment. *IEEE Transactions on Education*, 65(4):608–616, 2022.
- [4] M. Dileo, M. Zignani, and S. Gaito. Temporal graph learning for dynamic link prediction with text in online social networks. *Machine learning*, 113(4):2207–2226, 2024.
- [5] W. Fan, Y. Ma, Q. Li, Y. He, E. Zhao, J. Tang, and D. Yin. Graph neural networks for social recommendation. In *Proceedings of The World Wide Web Conference*, pages 417–426, 2019.
- [6] W. L. Hamilton, R. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1025–1035, 2017.
- [7] 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. In *Conference on Educational Data Mining (EDM)*, 2025.
- [8] M. M. Keikha, M. Rahgozar, and M. Asadpour. Deeplink: A novel link prediction framework based on deep learning. *Journal of Information Science*, 47(5):642–657, 2021.
- [9] S. Khoshraftar and A. An. A survey on graph representation learning methods. *ACM Transactions on Intelligent Systems and Technology*, 15(1):1–55, 2024.
- [10] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *Proceedings of the International Conference on Learning Representations*, 2017.
- [11] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [12] A. Krouska, C. Troussas, and M. Virvou. Social networks as a learning environment: Developed applications and comparative analysis. In *Proceedings of 8th International Conference on Information, Intelligence, Systems & Applications (IISA)*, pages 1–6, 2017.
- [13] M. Kumar, S. Mishra, S. S. Singh, and B. Biswas. Community-enhanced link prediction in dynamic networks. *ACM Transactions on the Web*, 18(2):1–32, 2024.
- [14] J. Li, H. Shomer, H. Mao, S. Zeng, Y. Ma, N. Shah, J. Tang, and D. Yin. Evaluating graph neural networks for link prediction: Current pitfalls and new benchmarking. *Advances in Neural Information Processing Systems*, 36:3853–3866, 2023.
- [15] D. Liben-Nowell and J. Kleinberg. The link prediction problem for social networks. In *Proceedings of the Twelfth International Conference on Information and Knowledge Management*, pages 556–559, 2003.
- [16] F. Mohamady and S. Dami. Link prediction in dynamic social networks using deep learning. In *2024 20th CSI International Symposium on Artificial Intelligence and Signal Processing (AISP)*, pages 1–5. IEEE, 2024.
- [17] H. Nassar, A. R. Benson, and D. F. Gleich. Pairwise link prediction. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 386–393, 2020.
- [18] A. Ninagawa and K. Eguchi. Link prediction using probabilistic group models of network structure. In *Proceedings of the 2010 ACM Symposium on Applied Computing*, pages 1115–1116, 2010.
- [19] V. Raghavan, G. Ver Steeg, A. Galstyan, and A. G. Tartakovsky. Modeling temporal activity patterns in dynamic social networks. *IEEE Transactions on Computational Social Systems*, 1(1):89–107, 2014.
- [20] R. Sahay, S. Nicoll, M. Zhang, T.-Y. Yang, C. Joe-Wong, K. A. Douglas, and C. G. Brinton. Predicting learning interactions in social learning networks: A deep learning enabled approach. *IEEE/ACM Transactions on Networking*, 31(5):2086–2100, 2023.
- [21] A. Sankar, Y. Liu, J. Yu, and N. Shah. Graph neural networks for friend ranking in large-scale social platforms. In *Proceedings of the Web Conference 2021*, pages 2535–2546, 2021.
- [22] M. Saqr and C. S. Montero. Learning and social networks - similarities, differences and impact. In *Proceedings of IEEE 20th International Conference on Advanced Learning Technologies (ICALT)*, pages 135–139, 2020.
- [23] K. Sharma, Y.-C. Lee, S. Nambi, A. Salian, S. Shah, S.-W. Kim, and S. Kumar. A survey of graph neural networks for social recommender systems. *ACM Computing Surveys*, 56(10):1–34, 2024.
- [24] J. Sosa, D. Martínez, and N. Guerrero. An unified approach to link prediction in collaboration networks. *arXiv preprint arXiv:2411.01066*, 2024.
- [25] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In *Proceedings of the International Conference on Learning Representations*, 2018.
- [26] C. Wang, V. Satuluri, and S. Parthasarathy. Local probabilistic models for link prediction. In *Seventh IEEE International Conference on Data Mining (ICDM 2007)*, pages 322–331, 2007.
- [27] W. Wei and K. M. Carley. Measuring temporal patterns

- in dynamic social networks. *ACM Transactions on Knowledge Discovery from Data*, 10(1):1–27, 2015.
- [28] S. Xu, C. Yang, C. Shi, Y. Fang, Y. Guo, T. Yang, L. Zhang, and M. Hu. Topic-aware heterogeneous graph neural network for link prediction. In *Proceedings of the 30th ACM international conference on information & knowledge management*, pages 2261–2270, 2021.
- [29] Z. Xu, H. Yuan, and Q. Liu. Student performance prediction based on blended learning. *IEEE Transactions on Education*, 64(1):66–73, 2021.
- [30] T.-Y. Yang, C. G. Brinton, and C. Joe-Wong. Predicting learner interactions in social learning networks. In *Proceedings of IEEE Conference on Computer Communications*, pages 1322–1330, 2018.
- [31] M. Zhang and Y. Chen. Weisfeiler-lehman neural machine for link prediction. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 575–583, 2017.
- [32] M. Zhang and Y. Chen. Link prediction based on graph neural networks. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 5171–5181, 2018.
- [33] J. Zhu. Max-margin nonparametric latent feature models for link prediction. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1179–1186, 2012.