

Expediting Support for Social Learning with Behavior Modeling

Yohan Jo, Gaurav Tomar, Oliver Ferschke, Carolyn P. Rosé,
Dragan Gašević

CMU-LTI-16-011

Language Technologies Institute
School of Computer Science
Carnegie Mellon University
5000 Forbes Ave., Pittsburgh, PA 15213
www.lti.cs.cmu.edu

Expediting Support for Social Learning with Behavior Modeling

Yohan Jo, Gaurav Tomar, Oliver Ferschke, Carolyn P. Rosé

SCHOOL OF COMPUTER SCIENCE
CARNEGIE MELLON UNIVERSITY
PITTSBURGH, PA, USA
{YOHANJ, GTOMAR, FERSCHKE, CPROSE}@CS.CMU.EDU

Dragan Gašević

SCHOOLS OF EDUCATION AND INFORMATICS
THE UNIVERSITY OF EDINBURGH
EDINBURGH, UK
DGASEVIC@ACM.ORG

Abstract

An important research problem for Educational Data Mining is to expedite the cycle of data leading to the analysis of student learning processes and the improvement of support for those processes. For this goal in the context of social interaction in learning, we propose a three-part pipeline that includes data infrastructure, learning process analysis with behavior modeling, and intervention for support. We also describe an application of the pipeline to data from a social learning platform to investigate appropriate goal-setting behavior as a qualification of role models. Students following appropriate goal setters persisted longer in the course, showed increased engagement in hands-on course activities, and were more likely to review previously covered materials as they continued through the course. To foster this beneficial social interaction among students, we propose a social recommender system and show potential for assisting students in interacting with qualified goal setters as role models. We discuss how this generalizable pipeline can be adapted for other support needs in online learning settings.

Chapter 1

Introduction

More and more recent work in educational data mining and learning analytics refers to a “virtuous cycle” of data leading to insight on what students need and then improvements in support for learning [21]. An important goal is tightening this cycle to improve learning experience. We are interested especially in social learning, drawing from a Vygotskian theoretical frame where learning practices begin within a social space and become internalized through social interaction. This may involve limited interaction, such as observation, or more intensive interaction through feedback, help exchange, sharing of resources, and discussion.

There are two main contributions of this paper. The first is to propose a pipeline and its component models that can expedite the cycle of data mining technology used to make sense of pathways of learner behaviors. The second is to present findings from an application of the proposed pipeline for the purpose of addressing a specific problem in goal-setting in a social learning platform.

Specifically, our first contribution is to propose a pipeline that can expedite the cycle of data infrastructure, learning process analysis, and intervention (Figure 1.1). Data infrastructure provides a uniform interface for heterogeneous data from social interaction in various platforms, such as connectivist Massive Open Online Courses (cMOOCs) [18], hobby communities, and Reddit communities, where people engage in follower-followee relations, post updates to their account, engage in threaded discussions, and also optionally link in blogs, YouTube videos, and other websites. Learning process analysis aims to analyze students’ processes depending on their social network configurations and to identify beneficial kinds of social connections. We developed a probabilistic graphical model that analyzes sequences of behaviors in terms of topics expressed and social media types that students actively engage in over time. Finally, intervention is introduced to foster beneficial social connections among students. We developed




Figure 1.1: Pipeline for educational data mining in social learning.

a recommender system that matches qualified students to discussions to increase opportunities for them to interact with other peers. The pipeline is iterative such that data from participation is used to create models that trigger interventions in subsequent runs of the course. Data from those later runs can be used to train new and better models in order to improve the interventions, and so on.

Our second contribution is to present findings from an application of the proposed pipeline to data from a social learning environment called ProSolo [15], in order to investigate the positive influence of observing goal-setting behavior. While goal-setting has been intensively researched and proven to be an important self-regulated learning (SRL) practice that often leads to success in learning, the influence of a student's goal-setting behavior on observers has little been investigated empirically. If goal-setting students turn out to be good role models, that is, beneficial to their social peers, we can encourage and help students to make such social connections with goal setters to enhance their learning experience. The usefulness of this effect may be especially desirable in online courses where the number of instructors is limited, or online communities that are not structured like courses, where students are required to take more agency in forging a learning path for themselves within an ecology of resources.

In the remainder of this paper, we begin by motivating the specifics of our pipeline as situated within the literature. Next, we present our pipeline and its application, along with concrete computational models and findings.

Chapter 2

Related Work

We first explore the literature on social learning and peer effects, which are the main context and motivation of our work. Next, we relate the components of our pipeline to related prior work.

2.1 Social Effects On Learning

Vygotsky’s view of social interaction as a key to learning and Bandura’s social learning theory [1] emphasize the importance of interaction to learning. In social contexts, by vicarious learning, students observe external models and learn from those observations even when not actively engaged in interaction [23]. Observation of role models facilitates motivation and self-efficacy for a task [17] and may be associated with positive changes in the observer’s behavior [12]. Drawing on this theoretical foundation, the positive impact of social interaction has been investigated in collaborative work [11] and in online courses [14]. Yet, to our knowledge, our work is the first to investigate goal-setting behavior specifically as a qualification of a role model in online learning.

2.2 Data Infrastructures

Several data infrastructures have been introduced to aid educational data mining for Massive Open Online Courses (MOOCs). For instance, MOOCdb [22] and DataStage¹, designed to store raw data from MOOCs, consolidate clickstream data from different MOOC platforms in a single, standardized database schema. This allows for developing platform-independent analysis tools, thus enabling analyses that span multiple courses hosted by different MOOC providers with reduced development effort. While these infrastructures focus on behavior data represented by clickstream logs, our proposed infrastructure deeply represents other aspects of student interactions, such as discussion behavior and social relationships, which require the natural language exchange between students.

DataShop [9] is another repository of learning data that focuses on the interaction between students and educational software. DataShop offers a set of tools for analyzing these datasets and building cognitive models that allow researchers to explore the relationships between students’

¹<http://datastage.stanford.edu/>

skills, concepts or misconceptions, and their trajectory in learning environments. Compared to human-computer interaction, however, analysis of human-human interaction is more complex, where interaction states are implicit, continuous, and controlled by a decentralized structure. Furthermore, in social learning, it is important to consider social relationships between humans and how those relationships moderate the effect of any computer agent or learning platform involved.

2.3 Learning Process Analysis

Analysis of students' learning processes has been a critical topic in education. Our method contributes to the literature on process mining through behavior modeling. Approaches to learning process analysis differ in the definition of the basic building block, often conceived of as states within a graph. Common building blocks for tutoring systems and educational games include knowledge components [26] and actions [16]. In dialogue settings, it is common to code each utterance according to a coding scheme and analyze the sequence of codes [5, 11]. In a MOOC context, states are often defined as course units [4, 8], course materials [4], and discussion threads [3].

Such predefined states, however, may not be the ideal units of states, especially in online courses where students can selectively engage in learning resources. Therefore, unsupervised modeling approaches are appealing for the purpose of identifying states that are meaningful indications of student interests obtained in a data-driven way. Markov models have been proposed to learn latent states and state transitions [19, 25]. However, their representation of a state is often too simple, e.g., a single multinomial distribution over observations. To improve the simple representation, the state transition topic model has been proposed [7], in which a state is represented as a mixture of topics, from which documents are generated via Latent Dirichlet Allocation (LDA). Yet, this type of model does not consider conditional state transitions, thereby imposing limitations in modeling the complex dynamics of social learning environments. Our model extends this basic approach so as to incorporate more information related to our interest, by distinguishing between multiple document types and conditioning state transitions on different types of social connection.

While we learn states based on the topics of student discussions in a course, there are also other views on the use of discussions for defining learning process. Milligan [10] argues that participation in discussion does not necessarily mean genuine learning of the content, so we should consider more latent and complex learning skills reflected in the content. Ezen-Can et al. [5] go deeper into discussion text to analyze the cognitive process evidenced in the discussion rather than the mere topics of the content.

2.4 Social Recommendation

In MOOCs, a student's learning process is affected by other peers especially through interaction in forums, which offer opportunities to develop communication and community. Hence, social recommendation algorithms can introduce appropriate students to certain discussions for

productive interaction. Suggested matches should be appropriate when viewed from both discussion and student sides [20], for example by suggesting a student to participate in discussions based on both the potential benefit of the student's expertise as an asset to the discussions while respecting the limitations of a student's resources for participation in more than a limited number of discussions [24]. Our model can recommend discussions to a student by balancing the benefit of the student's qualification to discussions, her relevance to discussions, and required effort.

Chapter 3

Three-Part Analytics Pipeline

Our pipeline is designed to expedite the process of exploiting student data leading to data-driven decision-making for enhancing student learning (Figure 1.1).

In this pipeline for social learning, the first component is a data infrastructure that maps diverse forms of social interaction into a common structure. This uniform interface allows the subsequent components—learning process analysis and intervention—to apply the same tools to different data, even from distinctly different discourse types, with little modification. Our development of this infrastructure, DiscourseDB¹, represents discourse-centered social interaction as an entity-relation model. Discourses (e.g., forums or social media) and individual contributions in a discourse (e.g., posts, comments, and utterances) are represented as generic containers generalizable to diverse social platforms. DiscourseDB also allows for defining arbitrary relations between contributions, e.g., a “reply-to” relation derived from the explicit reply structure of the platform versus one inferred through some automated analysis process. This flexibility helps the subsequent components of the pipeline avoid data-specific processing. DiscourseDB can store both active and passive activities of individuals, such as creating, revising, accessing, and following contributions, as well as forming social connections with other individuals. DiscourseDB is the key component of our pipeline, based on which the next components perform integrated analyses of discourses and social networking on multiple platforms with reusability.

The second component of our pipeline is analysis of students’ learning processes depending on their social connections. The goal is to assess students’ needs of support by understanding how learning processes are affected by social interaction and what types of social interaction are helpful to students. Just as Bayesian knowledge tracing enables modeling the learning process from a cognitive perspective and then supporting a student’s progress through a curriculum, Bayesian approaches can model learning processes at other levels, including supportive social processes. And similarly, these models can then be used to trigger support for the learning processes in productive ways. Hence, the third component of our pipeline draws upon insights obtained from the analysis to introduce interventions that can help students make beneficial social connections with other peers. We will propose two concrete examples of machine learning techniques for these two components in Section 5 and Section 6 respectively.

¹<http://discoursedb.github.io>

Chapter 4

Research Context

The remainder of the paper presents an example application of our general pipeline to a specific problem. This section describes the problem of our interest, data set, and terminology that will be used throughout the application.

4.1 Problem and Data

We examine goal-setting behavior as a potential qualification of good role models via learning process analysis and foster social connections with goal setters via recommendation support. Since most MOOCs and informal learning communities lack a measure to identify potentially good role models (e.g., a pretest), increased frequency of effective goal-setting behaviors may serve as an indirect indicator of success, as previous studies showed positive relationships between goal-setting behavior and learning outcomes [2, 6, 27].

The data was collected from an edX MOOC entitled *Data, Analytics, and Learning* (DAL-MOOC) [15], which ran from October to December 2014. This course covered theoretical principles about learning analytics as well as tutorials on social network analysis, text mining, and data visualization. This MOOC was termed a *dual layer* MOOC because students had the option of choosing a more standard path through the course within the edX platform or to follow a more self-regulated and social path in an external environment called ProSolo. The ProSolo layer allowed students to set their own learning goals and follow other students so that they could view activities and documents that offered clues about how to approach the course productively. While a huge literature on analysis of MOOC data focuses on Coursera, edX, and Udacity MOOCs, other platforms with more social affordances are growing in popularity. In order to serve the goal of identifying support needs and automating support that may be triggered in a social context, it is advantageous to work with data from socially-oriented platforms. We used the log data from ProSolo as our object of analysis, which include students' discussions on ProSolo and their own blogs and Twitter that they identified on their ProSolo profile pages, evidence of students' social connection with each other, and "goal notes," which students can use to set their learning goals in their own words.

We preprocessed discussion data before running our model. First, we filtered course-relevant tweets using the hashtags #prosolo, #dalmooc, and #learninganalytics. We confirmed that the

Goal notes	62	Tweets (relevant)	715
ProSolo posts	318	Tweets (irrelevant)	25,461
Blog posts	359		
Users	1,729	Social connections	814

Table 4.1: Descriptive statistics for ProSolo data.

tweets identified as irrelevant by this process have little to do with course activity. Because we are not interested in irrelevant content, we replaced such content with a tag to indicate irrelevant content. In order to prevent topics from being defined in terms of document types, we removed Twitter mentions and “RT” from tweets as well as other function words including URLs from all documents. Descriptive statistics for the data set are listed in Table 4.1.

4.2 Goal Quality and Social Connection

To categorize the quality of goal-setting behavior of each student, we first annotated each goal note written by students indicating whether it indeed contains a goal or not. 58% of goal notes contained goals. An example goal note is as follows: “*to understand learning analytics and see how these may be useful for my teaching and in particular, my learning resource design/development.*” On the basis of this annotation, we categorized students into three classes: (1) goal setters, (2) goal participants, and (3) goal bystanders. Goal setters have goal notes that mention their distal or/and proximal goals. Goal participants have goal notes, all of which are about something other than goals, e.g., experiences or questions. Goal bystanders have no goal notes. Note that the category of a student can change over time. All students start as goal bystanders and may become a goal participant or a goal setter as time passes. A student’s *social connection* is then categorized into seven classes: (S1) has already been following a goal setter, (S2) started to follow a goal setter at the current time point (S3) has been following a goal participant (but no goal setter), (S4) started to follow a goal participant at the current time point, (S5) has been following a goal bystander (at best), (S6) started to follow a goal bystander at the current time point, and (S7) follows no one. S2, S4, and S6 mean that a student’s social connection improved at the current time point, whereas S1, S3, and S5 indicate that a student remained in the same social connection category as in the previous time point.

Chapter 5

Learning Process Analysis

The second component of the pipeline aims to assess students' needs of support. Hence, we model students' behavior and analyze their learning processes especially as they experience changes in their social connections over time in the course. What models are best to use depends on the specific analysis of interest. In this section, we propose a Bayesian model designed for our problem and present our findings. In particular, we found out that a student's learning process is positively related to her social connections with goal setters. This pattern suggests the potential positive impact of supporting interaction with goal setters.

5.1 Model

Our model automatically extracts a representation of students' learning processes based on their discussions in a course and their social connections, which may reveal the influence of different configurations within the social space. We define the building blocks of learning processes, which we call states, in terms of discussed topics and the document types used for discussions (e.g. Twitter, blog). Given students' sequences of documents and social connection types over time, the model infers a set of meaningful states, along with the topics and document types for each state. The learned topics provide the information about students' interests, and the document types give an insight into how students use different media for different interests. The model also learns transition probabilities between states, conditioned on the category of a student's social connection in the source state.

We extend the state transition topic model from prior work [7] by including components for document types and conditional state transitions. More formally, our generative model assumes that each state has probability distributions over document types and over topics. Each state also has a probability distribution over states for each type of social connection, representing transition probabilities to states (including itself). This model assumes a generative process of students visiting states and making documents as follows. Suppose that there is a set of states that constitute students' learning processes. At each time point, a student enters into a state and writes documents; he chooses a document type and then repeats selecting a topic and writing a word. At the next time point, the student enters into a state according to the state transition probabilities of the current state, which is conditioned on the student's current social connection.

1. For each topic $j = 0, \dots, Z - 1$,
 - (a) Draw a word distribution $\phi_j \sim \text{Dirichlet}(\beta)$
2. For each state $c = 0, \dots, S - 1$,
 - (a) Draw a document type distribution $\psi_c \sim \text{Dirichlet}(\nu)$
 - (b) Draw a topic distribution $\theta_c \sim \text{Dirichlet}(\alpha)$
 - (c) For each category of social connection $b = 0, \dots, A$,
 - i. Draw a transition distribution $\pi_{cb} \sim \text{Dirichlet}(\gamma)$
3. For each time point $t = 1, 2, \dots$,
 - (a) Choose a state $s_t \sim \text{Categorical}(\pi_{s_{t-1}})$
 - (b) For each document,
 - i. Choose a document type $d \sim \text{Categorical}(\psi_{s_t})$
 - ii. For each word,
 - A. Choose a topic $z \sim \text{Categorical}(\theta_{s_t})$
 - B. Choose a word $w \sim \text{Categorical}(\phi_z)$

Figure 5.1: Generative process of our model.




Figure 5.2: Graphical representation of our model.

S	number of states
A	number of social category types
Z	number of topics
D	number of document types
M	number of sequences
docs_t^m	documents at t -th time point in sequence m
$w_{t,i}^m$	i -th word in docs_t^m
$z_{t,i}^m$	topic assigned to $w_{t,i}^m$
s_t^m, \bar{s}_t^m	state at t -th time point in sequence m
a_t^m	social status at t -th time point in sequence m
$\mathbf{w}_t^m, \bar{\mathbf{w}}_t^m$	all words in docs_t^m , $\mathbf{w} - \{\mathbf{w}_t^m\}$
$\mathbf{z}_t^m, \bar{\mathbf{z}}_t^m$	topics assigned to \mathbf{w}_t^m , $\mathbf{z} - \{\mathbf{z}_t^m\}$
\mathbf{d}_t^m	document types of docs_t^m
$N_{m,t,j}^{MTZ}$	number of words assigned topic j in docs_t^m
$N_{j,w}^{ZW}$	# words w assigned topic j
$N_{c,k}^{SD}$	# occurrences of document type d in state c
$N_{c,j}^{SZ}$	# words assigned topic j in state c
$N_{c,a,c'}^{SAS}$	# transitions from state c to c' given status a

Table 5.1: Description of notation.

This generative process is described more formally in Figure 5.1, and its graphical representation is shown in Figure 5.2. Notations are explained in Table 5.1.

Through an inference step, we can estimate the distributions over topics, document types, and transitions for each state (θ_c , ψ_c , and π_{cb} , respectively). So also would the topics (ϕ) and the state of each time point for students (s_t) be estimated. We use Gibbs sampling for inference. Each iteration samples $z_{t,i}^m$ and s_t^m according to the following probabilities.

$$\begin{aligned}
p(z_{t,i}^m = j | \bar{\mathbf{z}}_{t,i}^m, \mathbf{w}, s_t^m) &\propto (N_{m,t,j}^{MTZ} + \alpha) \frac{N_{j,w_{t,i}^m}^{ZW} + \beta}{\sum_{w'} (N_{j,w'}^{ZW} + \beta)}, \\
p(s_t^m = c | \bar{\mathbf{s}}_t^m, \mathbf{z}, \mathbf{a}, \mathbf{d}) &\propto \left(\prod_{k=1}^D \frac{\Gamma(N_{c,k}^{SD} + \nu + N_{m,t,k}^{MTD})}{\Gamma(N_{c,k}^{SD} + \nu)} \frac{\Gamma(\sum_{k'} (N_{c,k'}^{SD} + \nu))}{\Gamma(\sum_{k'} (N_{c,k'}^{SD} + \nu) + |\mathbf{d}_t^m|)} \right) \\
&\times \left(\prod_{j=1}^Z \frac{\Gamma(N_{c,j}^{SZ} + \alpha + N_{m,t,j}^{MTZ})}{\Gamma(N_{c,j}^{SZ} + \alpha)} \frac{\Gamma(\sum_{j'} (N_{c,j'}^{SZ} + \alpha))}{\Gamma(\sum_{j'} (N_{c,j'}^{SZ} + \alpha) + |\mathbf{z}_t^m|)} \right) \\
&\times \left(\frac{N_{s_{t-1}^m, a_{t-1}^m, c}^{SAS} + \gamma}{\sum_{c'} (N_{s_{t-1}^m, a_{t-1}^m, c'}^{SAS} + \gamma)} \right) \\
&\times \left(\frac{N_{c, a_t^m, s_{t+1}^m}^{SAS} + \mathbf{1}(s_{t-1}^m = c = s_{t+1}^m) + \gamma}{\sum_{c'} (N_{c', a_t^m, s_{t+1}^m}^{SAS} + \mathbf{1}(s_{t-1}^m = c' = s_{t+1}^m) + \gamma)} \right).
\end{aligned}$$

From the sampling results, we can estimate

$$\begin{aligned}
\phi_{j,w} &= \frac{N_{j,w}^{ZW} + \beta}{\sum_{w'} (N_{j,w'}^{ZW} + \beta)}, \theta_{c,j} = \frac{N_{c,j}^{SZ} + \alpha}{\sum_{j'} (N_{c,j'}^{SZ} + \alpha)}, \\
\psi_{c,k} &= \frac{N_{c,k}^{SD} + \nu}{\sum_{k'} (N_{c,k'}^{SD} + \nu)}, \theta_{t,j}^m = \frac{N_{m,t,j}^{MTZ} + \alpha}{\sum_{j'} (N_{m,t,j'}^{MTZ} + \alpha)}, \\
\pi_{c,b,c'} &= \frac{N_{c,b,c'}^{SAS} + \gamma}{\sum_{c'} (N_{c,a,c'}^{SAS} + \gamma)},
\end{aligned}$$

where θ_t^m is the topic distribution of docs $_t^m$. A detailed derivation process and the source code are available on our website¹. We can also infer the state of each time point of a student, e.g., by the state assigned to each time point during the sampling process. For an unseen sequence of documents, we may infer the state of each time point using a Viterbi algorithm based on the document type distribution, topic distribution, and state transition distribution of each state. Once states are finalized, maximum likelihood estimation can be applied to estimate the topic distribution of the documents in each time point.

¹<http://cs.cmu.edu/~yohanj/research>

State	Topics	RelGoalNote	IrGoalNote	Post	Blog	RelTweet	IrTweet
0	Course-irrelevant tweets	0.00	0.00	0.00	0.00	0.00	1.00
1	Concept map, network analysis (Week 9)	0.00	0.00	0.02	0.01	0.18	0.78
2	Social capital (Week 3)	0.04	0.01	0.19	0.30	0.18	0.27
3	Tableau (Week 2), Gephi (Week 3), Lightside (Week 7)	0.01	0.03	0.10	0.28	0.24	0.34
4	Prediction models (Week 5)	0.01	0.02	0.29	0.22	0.10	0.36
5	Data wrangling (Week 2)	0.01	0.01	0.12	0.08	0.26	0.52
6	Visualization (Week 3)	0.05	0.02	0.24	0.47	0.08	0.15
7	Epistemology, assessment, pedagogy (Week 4)	0.05	0.00	0.18	0.22	0.30	0.25
8	Prediction, decision trees (Week 5)	0.02	0.02	0.19	0.40	0.09	0.28
9	Share, creativity (mixed topics)	0.00	0.02	0.12	0.13	0.21	0.52

Table 5.2: Learned states with their topics and document type distribution (each row sums to 1). (RelGoalNote: goal notes containing a goal, IrGoalNote: goal notes without a goal, Post: posts on ProSolo, Blog: personal blog posts, RelTweet: course-relevant tweets, IrTweet: course-irrelevant tweets)

5.2 Findings

We applied the model to the ProSolo data and examined the correlation between the categories of social connection and learning behaviors. We ran our model with the number of states set to 10 and the number of topics set to 20. We defined the unit of a time point as one week, and if a student had no activity in a certain week, that week was omitted from her sequence.

5.2.1 Learned States

This section examines the states learned by the model and whether they align with course units or suggest different behavior units. Table 5.2 summarizes the learned states with their topics and document type distributions (interpreted from ϕ , θ_c , and ψ_c). Most states are aligned with course units covering important course topics, such as learning analytics, data visualization, social networks, and prediction models. However, State 0 is where students do not participate in course discussion but post course-irrelevant tweets. State 3 is about hands-on practice of software tools across the course, and State 9 covers many side topics.

Document types and their correlations with topics in each state also reveal interesting student behaviors (Table 5.2). Tweets tend to take a large proportion and goal notes a small proportion in every state due to their relative volumes. According to our examination of the data, blog posts are actively used for summarizing readings and tutorials, and tweets are used as a means of

	Social Connection			
	GS _{S1+S2}	GP _{S3+S4}	GB _{S5+S6}	NO _{S7}
# Time Points	139	315	265	821
% Time Points				
State 0	0.59**	0.75	0.75	0.71
State 1	0.17*	0.10	0.03	0.04
State 2	0.05	0.02	0.02	0.04
State 3	0.04*	0.00	0.01	0.01
State 4	0.01	0.02	0.03	0.06
State 5	0.05	0.03	0.06	0.05
State 6	0.05	0.02	0.02	0.02
State 7	0.03	0.01	0.03	0.02
State 8	0.00	0.03	0.02	0.02
State 9	0.01	0.04	0.03	0.04

Table 5.3: Proportion of time points students stay in each state depending on the social connection (each column sums to 1). “**” and “*” indicate that GS is significantly different from other categories in bold with $p < 0.01$ and $p < 0.05$, respectively, by Pearson’s chi-square test. GS, GP, and GB each represent either “has been following” or “started to follow” a goal setter, a goal participant, and a goal bystander, respectively. NO means to follow no one.

communicating with lecturers (e.g., State 5). ProSolo posts are most accessible to ProSolo users, so students use them to reveal their opinions and questions.

5.2.2 Students Following Goal Setters

We investigated the learning processes of the students who follow goal setters and their positive learning behaviors, based on the number of weeks students spent in each state (Table 5.3) and state transition patterns (Figure 5.3). Note that in Table 5.3, for simplicity, “has been following” and “started to follow” categories are combined in each column.

Twitter usage: Students generally wrote irrelevant tweets (State 0) significantly more than course-related documents regardless of the social category. However, the students following goal setters spent noticeably fewer weeks in this state.

Participation duration: The topics of the states in which students stay reveal how long they persist in the course. The students following goal setters were more likely to discuss the material taught in the last week (State 1) than other students, that is, they were active in the last phase of the course. This phenomenon is consistent with the previous finding that social connection can lower attrition rates in MOOCs [14].

Activities of interest: The number of weeks students spend in each state reflects the activities students are interested in. The students following goal setters were more active in hands-on practice (State 3) than other students. This state captures active learning, i.e., hands-on practice of software tools such as Tableau, Gephi, and Lightside across the course. Hands-on practice

requires higher motivation than merely watching lectures, so these students might have been helped by observation of role models as discussed in the literature [17]. This trend would have not been as clear using predefined states based on course units [4], which could not learn to distinguish lecture and practice activities.

Study habits or challenges: Transition patterns may reveal students' study habits or challenges. Figure 5.3a shows frequent transitions between three states (States 1, 3, and 5) that are associated with materials taught in different weeks. Such transitions may reflect the SRL strategy of activating and applying prior knowledge to the current situation [13].

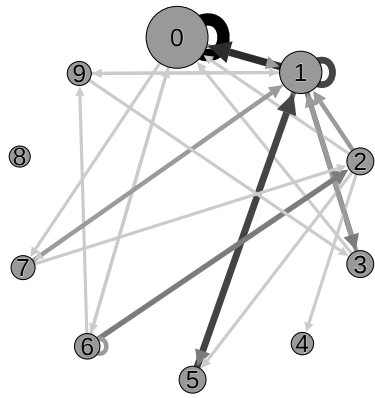
These positive effects associated with following goal setters are not apparent with other social connection types, such as following goal participants or goal bystanders. The students who start to follow a goal setter (Figure 5.3b) begin to show the behaviors of those following a goal setter (Figure 5.3a) changing from the behaviors of those following a goal participant (Figure 5.3c). This indicates that "who to follow" is more important than simply following someone.

5.2.3 Students With No Social Connection

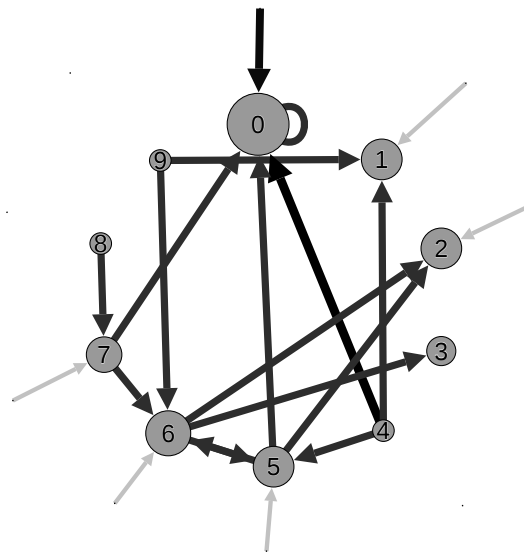
Students with no social connection were not passive users of ProSolo. Compared to students following goal setters, however, they were far more passive in writing documents and engaging with discussions. They also showed the following learning behavior (Table 5.3).

Twitter usage: Interestingly, the relative amount of time they spent in State 0 is similar to students who had social connections. They may be passive in using social media, or they may have not bothered to reveal their Twitter information.

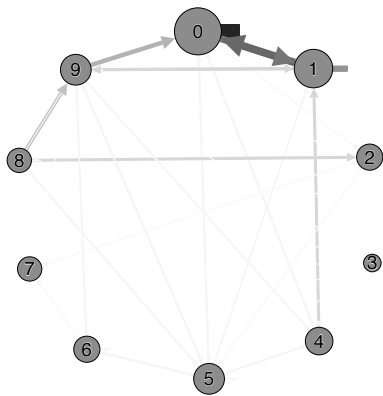
Activities of interest: While students with social connections had preferred states, these students were spread across all states quite evenly. They may have diverse interests but be indifferent about making connections with other people.



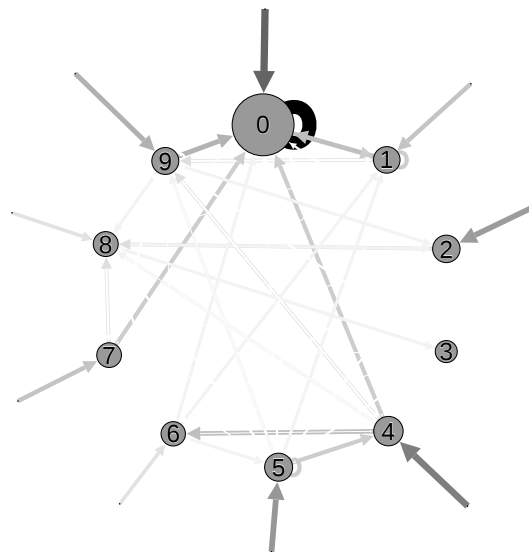
(a) S1. Has been following a goal setter



(b) S2. Started to follow a goal setter



(c) S3. Has been following a goal participant



(d) S7. Follows no one

Figure 5.3: State transition patterns. Nodes are states whose size reflects the number of weeks students visit the states. Edges are transitions whose thickness and darkness reflect transition frequency. Edges without a source node represent the probability of being the first state in a learning path.

Chapter 6

Intervention for Support

On the basis of the insights obtained from the previous component, the third component of our pipeline is to offer appropriate support, especially towards fostering beneficial social connections between students. We argue that a recommender system can serve this purpose, by presenting its potential positive impact as assessed on the corpus.

6.1 Model

Our recommender system aims to match qualified students (e.g., goal setters) to discussions so that they can interact with and benefit the discussants through discussions. Our model has two steps: relevance prediction and constraint filtering. The relevance prediction step learns the relevance between students and discussions using student- and discussion-related features that are potentially valuable in making recommendations. The learned relevance reflects students' preferences and tendencies, but may not reflect the ideal matches for fostering learning. The constraint filtering step thus combines the relevance scores with some constraints that foster interaction between qualified students and other students, and finalizes recommendations.

6.1.1 Relevance Prediction

Our algorithm extends the earlier model proposed by Yang et al. [24], which is also designed to match students with discussions, by incorporating additional components related to students' qualifications. The relevance matrix between students and discussions is denoted as $R = \{r_{u,d}\}$ for every student u and discussion d . $r_{u,d}$ is 1 if and only if student u has participated in discussion d . As in the original model, our model exploits student features, discussion features, and implicit feedback. In addition, we add the following additional student features related to qualifications of interest.

- **Goal quality (λ):** A student's goal quality as defined in Section 4. (2: goal setter, 1: goal participant, 0: goal bystander)
- **Degree centrality (ψ):** The average of the authority and hub scores of a student's network. High centrality may serve as a hub for further social interaction.

Taking these additional features into account, our new relevance prediction model can be formulated as:

$$r_{u,d} = bias + (P_u + \phi_u \Phi + \theta_u \Theta + \lambda_u \Lambda + \psi_u \Psi + \Gamma_\gamma)^T \times (Q_d + \delta_d \Delta + l_d L + \frac{1}{\sqrt{|U(d)|}} \sum_{v \in U(d)} \varphi_v). \quad (6.1)$$

ϕ_u and θ_u are the number of discussions u has participated in and the number of discussions u has initiated. Γ_γ is a one-hot vector indicating the course week that u registered for the course; this may be related to the student's motivation. δ_d and l_d are the number of replies and the length of the content in d . $U(d)$ denotes the set of students participating in d , and φ_v is the predicted preference v of u . P_u and Q_d are the biases of u and d . Φ , Θ , Λ , Ψ , Δ , and L are one-dimensional feature weights. Additional details regarding parameter estimation are available in Yang et al.'s paper [24].

6.1.2 Constraint Filtering

It is important for students to receive relevant discussion recommendations, but it is also desirable to recommend students to discussions where the other discussants can benefit from the student's participation and thereby to optimize the overall community welfare. For this purpose, we use a max cost flow model to subject the relevance scores obtained in the previous step with the following constraints.

1. **Goal quality:** For every discussion d , at least one student u to which we recommend d should have a goal quality G_u greater than some threshold G .
2. **Degree centrality:** For every discussion d , at least one student u to which we recommend d should have a centrality score C_u greater than some threshold C . The rationale behind this constraint is that students with high centrality may serve as a hub through which more students are connected.
3. **Workload:** Qualified students should not be matched to too many discussions so that their workload is minimized.

These constraints are formulated into the following optimization problem, where our goal is to compute $f_{u,d} \in \{0, 1\}$, an indicator of whether discussion d is recommended to student u , that maximizes the objective function:

$$\begin{aligned} \max \quad & \sum_{u,d} f_{u,d} \cdot r_{u,d} - \alpha \cdot \sum_d \sum_u \mathbb{1}(G_u \cdot f_{u,d} \geq G)(G_u - G) \\ & - \alpha \cdot \sum_d \sum_u \mathbb{1}(C_u \cdot f_{u,d} \geq C)(C_u - C) \quad \text{s.t.} \\ & \forall d \in D, \exists u \in U, G_u \cdot f_{u,d} \geq G, \\ & \forall d \in D, \exists u \in U, C_u \cdot f_{u,d} \geq C, \end{aligned} \quad (6.2)$$

where D and U are the sets of discussions and students, respectively. $\mathbb{1}$ is an indicator function. The second and third terms prevent qualified students from being assigned to too many

discussions. The concept of a concave cost network for solving this optimization problem is detailed in Yang et al.’s paper [24].

6.2 Findings

Since we have identified positive learning behaviors of students who follow goal setters, we may want to support students by fostering interaction with goal setters. Instead of recommending direct following relations, which are not supported by many learning platforms, we recommend discussions to qualified students so that they can interact with the discussants. We first assess the extent to which students are sensitive to qualified students prior to explicit intervention, and then present the potential added value of our recommendation model.

6.2.1 Students’ Awareness of Role Models

Our first step is to assess whether students can identify effective role models in discussion activities (ProSolo posts), by measuring the impact of the information about students’ qualifications on the prediction of discussion participation. This task is to infer links between students and discussions that we hid from an observed static snapshot of a network of discussion participation based on observable data. A measured positive impact here would indicate some sensitivity on the part of students to interact with qualified students naturally. We train a predictive model of students’ participation in discussions on two thirds of student-discussion pairs. We then predict the discussion participation of the remaining pairs. Our evaluation metric is mean average precision (MAP).

We compared four configurations by varying the information about students’ qualifications that is used as feature for relevance prediction. In particular, CAMF uses only basic features, such as the numbers of discussions each student initiated and participated in and each discussion’s length, number of replies, and participants. CAMF_G and CAMF_C add information about goal quality and degree centrality, respectively, and CAMF_GC adds both. The evaluation was conducted as a link prediction task, based on the relevance scores predicted in the relevance prediction step. Students’ qualification information did not improve link prediction accuracy (Table 6.1). This means that students are not proactively sensitive to peers’ qualifications while participating in discussions, which supports our view that explicit recommendation could be valuable for encouraging students to interact with qualified peers through discussions.

6.2.2 Recommendation Quality

The recommendation of discussions should be consistent with both the relevance between students and discussions (the relevance prediction step) and constraints for beneficial social connection (the constraint filtering step). To this end, we evaluated recommendation quality on Overall Community Benefit (OB) (Equation 6.2 without the constraints): the relevance of our recommendations penalized by the burden on the students induced by the recommendations. The higher OB the better.

Configuration	MAP	Configuration	MAP
CAMF	0.465	CAMF_C	0.455
CAMF_G	0.438	CAMF_GC	0.439

Table 6.1: MAP for link prediction.

Configuration	OB	Configuration	OB
GoalPart	1.888	MCCF_G	3.683
HighCent	1.943	MCCF_C	3.770
GoalPart_HighCent	1.873	MCCF_GC	3.656

Table 6.2: Overall Community Benefit for recommendation.

We tested three configurations by varying the constraints incorporated into the constraint filtering step. **MCCF_G** requires that every discussion have at least one goal participant or goal setter. **MCCF_C** requires that every discussion have at least one student whose degree centrality is higher than 0.1. **MCCF_GC** requires both. In addition, the following configurations were tested as baseline without incorporation into the model. **GoalPart** filters goal participants or goal setters after making recommendations based on predicted relevance. Similarly, **HighCent** filters students with degree centrality higher than 0.1. **GoalPart_HighCent** filters goal participants or goal setters with degree centrality higher than 0.1. Incorporating the constraints about students' goal quality and degree centrality into the model (**MCCF_G**, **MCCF_C**, and **MCCF_GC**) achieved higher OB than the simple filtering approaches (Table 6.2). That is, our algorithm effectively matches qualified models to relevant discussions in such a way that students in every discussion can interact with qualified models while balancing the load of the models.

Chapter 7

Discussion

According to our learning process analysis, students benefit from social connections with effective goal setters through ProSolo’s follower-followee functionality. They stay longer in the course, engage in hands-on practices, and link materials across the course. This supports the view that goal-setting behavior is a useful qualification for potential role models. According to the discussion participation prediction task, explicit intervention is important for helping students be aware of qualified students and interact with them via discussions. Therefore, we incorporated the information about students’ qualifications into our recommendation model as constraints, successfully matching qualified learning partners to relevant discussions.

This work started from the need for expediting data analysis and analysis-informed support in social learning where students interact with one another via various social media in order to pursue their own learning goals. This expedition builds on DiscourseDB, data infrastructure for complex interaction data from heterogeneous platforms. We proposed a probabilistic graphical model to analyze students’ learning processes depending on the state of their social connections, and proposed a recommender system that can improve student support on the basis of the insights obtained from the analysis. This pipeline arguably should allow us to apply the techniques to different learning communities with little effort.

Goal-setting behavior is an important practice in SRL and is known to be difficult for students, so an analysis towards improvement of this skill is arguably valuable. Nevertheless, in this study we have not examined how this behavior influences the domain learning of students. This is due both to the limited data size for our first trial to use ProSolo in MOOCs as well as a lack of learning gain measures. However, the modeling techniques proposed in this paper can readily be applied to other data sets if the requisite data become available. We are also interested in investigating different SRL strategies besides goal-setting in social learning, and how social interaction influences the SRL behaviors of the students. Ultimately, the real value of the work will be demonstrated not with a corpus analysis, as for our proposed recommendation approach, but with an intervention study in a real MOOC. We are working towards incorporating this approach in a planned rerun of DALMOOC.

Acknowledgments

This research was supported by the National Science Foundation under grants ACI-1443068 and IIS-1320064, and by the Naval Research Laboratory and Google.

Bibliography

- [1] Albert Bandura. *Social Learning Theory*. Morristown, N. J.: General Learning Press, 1971. 2.1
- [2] Ann Battle and Allan Wigfield. College women’s value orientations toward family, career, and graduate school. *Journal of Vocational Behavior*, 62(1):56–75, 2003. ISSN 00018791. doi: 10.1016/S0001-8791(02)00037-4. URL <http://www.sciencedirect.com/science/article/pii/S0001879102000374>. 4.1
- [3] Alejandro Bogarín and Rebeca Cerezo. Discovering students’ navigation paths in Moodle. In *EDM ’15*, pages 556–557, 2015. 2.3
- [4] Carleton Coffrin, Linda Corrin, Paula de Barba, and Gregor Kennedy. Visualizing patterns of student engagement and performance in MOOCs. In *LAK ’14*, pages 83–92, March 2014. ISBN 9781450326643. doi: 10.1145/2567574.2567586. URL <http://dl.acm.org/citation.cfm?id=2567574.2567586>. 2.3, 5.2.2
- [5] Aysu Ezen-Can, Kristy Elizabeth Boyer, Shaun Kellogg, and Sherry Booth. Unsupervised modeling for understanding MOOC discussion forums. *LAK ’15*, pages 146–150, 2015. doi: 10.1145/2723576.2723589. URL <http://dl.acm.org/citation.cfm?id=2723576.2723589>. 2.3
- [6] Jenefer Husman and Duane F. Shell. Beliefs and perceptions about the future: A measurement of future time perspective. *Learning and Individual Differences*, 18(2):166–175, April 2008. ISSN 10416080. doi: 10.1016/j.lindif.2007.08.001. URL <http://www.sciencedirect.com/science/article/pii/S1041608007001148>. 4.1
- [7] Yohan Jo and Carolyn Penstein Rosé. Time Series Analysis of Nursing Notes for Mortality Prediction via a State Transition Topic Model. In *CIKM ’15*, 2015. ISBN 9781450337946. 2.3, 5.1
- [8] René F. Kizilcec, Chris Piech, and Emily Schneider. Deconstructing disengagement: Analyzing Learner Subpopulations in Massive Open Online Courses. In *LAK ’13*, page 170, April 2013. ISBN 9781450317856. doi: 10.1145/2460296.2460330. URL <http://dl.acm.org/citation.cfm?id=2460296.2460330>. 2.3
- [9] Kenneth R. Koedinger, John C. Stamper, Brett Leber, and Alida Skogsholm. Learnlab’s datashop: A data repository and analytics tool set for cognitive science. *Topics in Cognitive Science*, 5(3):668–669, 2013. ISSN 1756-8765. doi: 10.1111/tops.12035. URL <http://dx.doi.org/10.1111/tops.12035>. 2.2
- [10] Sandra Milligan. Crowd-sourced learning in MOOCs. In *LAK ’15*, pages 151–155, March

2015. ISBN 9781450334174. doi: 10.1145/2723576.2723596. URL <http://dl.acm.org/citation.cfm?id=2723576.2723596>. 2.3

- [11] Inge Molenaar and Ming Ming Chiu. Effects of sequences of socially regulated learning on group performance. In *LAK '15*, pages 236–240, March 2015. ISBN 9781450334174. doi: 10.1145/2723576.2723586. URL <http://dl.acm.org/citation.cfm?id=2723576.2723586>. 2.1, 2.3
- [12] Elizabeth Levy Paluck, Hana Shepherd, and Peter M Aronow. Changing climates of conflict: A social network experiment in 56 schools. *Proceedings of the National Academy of Sciences*, 113(3):566–571, January 2016. 2.1
- [13] Paul R. Pintrich. A Conceptual Framework for Assessing Motivation and Self-Regulated Learning in College Students. *Educational Psychology Review*, 16(4):385–407, December 2004. ISSN 1040-726X. doi: 10.1007/s10648-004-0006-x. URL <http://link.springer.com/10.1007/s10648-004-0006-x>. 5.2.2
- [14] Carolyn Penstein Rosé, Ryan Carlson, Diyi Yang, Miaomiao Wen, Lauren Resnick, Pam Goldman, and Jennifer Sherer. Social factors that contribute to attrition in MOOCs. In *L@S '14*, pages 197–198, March 2014. ISBN 9781450326698. doi: 10.1145/2556325.2567879. URL <http://dl.acm.org/citation.cfm?id=2556325.2567879>. 2.1, 5.2.2
- [15] Carolyn Penstein Rosé, Oliver Ferschke, Gaurav Tomar, Diyi Yang, Iris Howley, Vincent Aleven, George Siemens, Matthew Crosslin, Dragan Gasevic, and Ryan Baker. Challenges and Opportunities of Dual-Layer MOOCs: Reflections from an edX Deployment Study. In *CSCL '15*, pages 848–851, 2015. 1, 4.1
- [16] Elizabeth Rowe, Ryan S. Baker, and Jodi Asbell-Clarke. Strategic game moves mediate implicit science learning. In *EDM '15*, pages 432–436, 2015. 2.3
- [17] Dale H Schunk and Antoinette R Hanson. Peer models: Influence on children’s self-efficacy and achievement. *Journal of educational psychology*, 77(3):313–322, 1985. ISSN 0022-0663. doi: 10.1037/0022-0663.77.3.313. 2.1, 5.2.2
- [18] George Siemens. Connectivism: A learning theory for the digital age. *International Journal of Instructional Technology and Distance Learning*, 2014. 1
- [19] Kevin Tang, Li Fei-Fei, and Daphne Koller. Learning latent temporal structure for complex event detection. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1250–1257, June 2012. ISBN 978-1-4673-1228-8. doi: 10.1109/CVPR.2012.6247808. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6247808>. 2.3
- [20] Loren Terveen and David W McDonald. Social matching: A framework and research agenda. *ACM transactions on computer-human interaction*, 12(3):401–434, 2005. 2.4
- [21] Candace Thille. Education Technology as a Transformational Innovation. *White House Summit on Community Colleges: Conference Papers*, pages 73–78, 2010. 1
- [22] Kalyan Veeramachaneni, Sherif Halawa, Franck Dernoncourt, Una-May O’Reilly, Colin Taylor, and Chuong Do. Moadb: Developing standards and systems to support MOOC data science. *CoRR*, abs/1406.2015, 2014. URL <http://arxiv.org/abs/1406>.

2015. 2.2

- [23] P. H. Winne and a. F. Hadwin. Self-regulated learning and socio-cognitive theory. *International Encyclopedia of Education*, pages 503–508, 2010. doi: 10.1016/B978-0-08-044894-7.00470-X. 2.1
- [24] Diyi Yang, David Adamson, and Carolyn Penstein Rosé. Question recommendation with constraints for massive open online courses. In *RecSys '14*, pages 49–56, 2014. 2.4, 6.1.1, 6.1.1, 6.1.2
- [25] Jaewon Yang, Julian McAuley, Jure Leskovec, Paea LePendou, Nigam Shah, and Biomedical Informatics. Finding Progression Stages in Time-evolving Event Sequences. In *WWW '14*, pages 783–793, April 2014. ISBN 9781450327442. doi: 10.1145/2566486.2568044. URL <http://dl.acm.org/citation.cfm?id=2566486.2568044>. 2.3
- [26] Chengling Zhao and Liyong Wan. A shortest learning path selection algorithm in e-learning. *Int'l Conference on Advanced Learning Technologies*, pages 94–95, 2006. doi: 10.1109/ICALT.2006.1652375. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1652375>. 2.3
- [27] Barry J. Zimmerman. Goal setting: A key proactive source of academic self-regulation. In *Motivation and self-regulated learning: Theory, research, and applications*, pages 267–295. Erlbaum, 2008. 4.1