

Interpretable Engagement Models for MOOCs using Hinge-loss Markov Random Fields

Arti Ramesh¹, Dan Goldwasser², Bert Huang³, Hal Daume III⁴, and Lise Getoor⁵

SUNY Binghamton¹ Purdue University² Virginia Tech³

University of Maryland, College Park⁴ University of California, Santa Cruz⁵

artir@binghamton.edu, dgoldwas@purdue.edu, bhuang@vt.edu, hal@cs.umd.edu, getoor@ucsc.edu

Abstract—Maintaining and cultivating student engagement is critical for learning. Understanding factors affecting student engagement can help in designing better courses and improving student retention. The large number of participants in massive open online courses (MOOCs) and data collected from their interactions on the MOOC open up avenues for studying student engagement at scale. In this work, we develop an interpretable statistical relational learning model for understanding student engagement in online courses using a complex combination of behavioral, linguistic, structural, and temporal cues. We show how to abstract student engagement types of active, passive, and disengagement as meaningful latent variables using logical rules in our model connecting student behavioral signals with student success in MOOCs. We demonstrate that the latent formulation for engagement helps in predicting two measures of student success: performance, their final grade in the course, and survival, their continued presence in the course till the end, across seven MOOCs. Further, in order to initiate better instructor interventions, we need to be able to predict student success early in the course. We demonstrate that we can predict student success early in the course reliably using the latent model. We also demonstrate the utility of our models in predicting student success in new courses, by training our models on one course and testing on another course. We show that the latent abstractions are helpful in predicting student success and engagement reliably in new MOOCs that haven't yet gathered student interaction data. We then perform a closer quantitative analysis of different features derived from student interactions on the MOOC and identify student activities that are good indicators of student success at different points in the course. Through a qualitative analysis of the latent engagement variable values, we demonstrate their utility in understanding students' engagement levels at various points in the course and movement of students across different types of engagement.

Index Terms—latent engagement models, student engagement, graphical models, statistical relational models, course success prediction

1 INTRODUCTION

The large number of students participating in MOOCs provides the opportunity to perform rich analysis of large-scale online interaction and behavioral data. This analysis can help improve student engagement in MOOCs by identifying patterns, suggesting new feedback mechanisms, and guiding instructor interventions. Additionally, insights gained by analyzing online student engagement can also help validate and refine our understanding of engagement in traditional classrooms.

In this work, we study the different aspects of online student behavior in MOOCs, develop a large-scale, data-driven approach for modeling student engagement. We study two course *success indicators* for online courses—1) *performance*: how well the student performs in the graded elements in the courses, and 2) *survival*: whether the student follows the course to completion. We demonstrate the construction of a holistic model incorporating content (e.g., language), structure (e.g., social interactions in discussion forums), and outcome data and show that jointly measuring different aspects of student behavior early in the course can provide a strong indication of course success indicators.

Examining real MOOC data, we observe that there are several indicators useful for gauging students' engagement, such as viewing course content, interacting with other students or instructors on the discussion forums, and the topic and tone of these interactions. Furthermore, students often engage in different

aspects of the course throughout its duration. For example, some students engage in the social aspects of the online community—by posting in forums and asking and answering questions—while others only watch lectures and take quizzes without interacting with the community. We take these differences into account and propose a model that uses the different behavioral aspects to distinguish between forms of engagement: passive, active, and disengagement. We use these engagement types to predict student success, and reason about their behavior over time.

Predictive modeling over MOOC data poses a significant technical challenge requiring the ability to combine language analysis of forum posts with graph analysis over very large networks of entities (students, instructors, assignments, etc.). To address this challenge, we use a recently developed statistical relational learning framework—*hinge-loss Markov random fields* (HL-MRFs). This framework provides an easy means to represent and combine behavioral, linguistic, and structural features in a concise manner. Our model is specified using weighted first-order logic rules, thus making it easy to encode and interpret how different behavioral, linguistic, structural, and temporal signals are indicative of different types of engagement and student success. Our first contribution is constructing a holistic model to represent and reason about various student activities in the MOOC setting. Our work is a step toward helping educators understand how students interact on MOOCs.

Our second contribution is providing a data-driven formulation

that captures student engagement in the MOOC setting. As in the traditional classroom setting, assessing online student engagement requires interpretation of indirect cues. Identifying these cues in an electronic setting is challenging, but the large amounts of available data can offset the loss of in-person communication. We analyze students' online behavior to identify how they engage with course materials and investigate how engagement can be helpful in predicting student performance and survival in the course. We extend our HL-MRF model to encode engagement as *latent* variables, which take into account the observed behaviors of online students and their resulting performance and survival in the class. The latent engagement variables in our model represent three prominent forms of engagement: 1) active engagement, 2) passive engagement, and 3) disengagement. Uncovering these different latent engagement states for students provides a better explanation of students' behavior leading to course completion and resulting grades.

We apply our models to real data collected from seven Coursera¹ courses at University of Maryland, College Park and empirically show their ability to capture behavioral patterns of students and predict student success. Our experiments validate the importance of providing a holistic view of students' activities, combining all aspects of online behavior, in order to accurately predict the students' motivation and ability to succeed in the class. We conduct experiments to evaluate two important course success parameters in online courses: course performance and survival. Early detection of changes in student engagement can help educators design interventions and adapt the course presentation to motivate students to continue with the course [1]. We show that our model is able to make meaningful predictions using data obtained at an *early* stage in the class. These predictions can help provide a basis for instructor intervention at an early stage in the course, helping to improve student retention rates. Further, we evaluate the strength of our models in predicting student survival on unseen courses and demonstrate that our models are able to make meaningful predictions for previously unseen courses, even at an early stage in the course. We also perform a comprehensive feature evaluation in predicting student success in MOOCs in different time periods of the course. Our interpretable probabilistic framework helps in encoding the different feature dependencies and evaluating their individual and combined effect on student success and engagement. Our findings strengthen the importance of using a holistic model and uncover important details about student interactions that is helpful for instructors. Finally, we use the latent engagement variables to unearth patterns in student engagement over the course of the class and detect changes in engagement. This can be potentially used by instructors to understand student movement from one engagement type to another and initiate interventions.

This work expands on the work described in [2], by providing additional experimental results. We look into several measures of student success, such as predicting student performance, predicting final student survival, and early prediction of student survival, building on our work in [3] and [2], and provide experimental results for seven MOOCs, covering a wide range of topics. We also include a suite of results for predicting student survival, predicting student survival at early time periods, predicting student survival for unseen courses, and predicting student survival early for unseen courses. We also include a comprehensive analysis

of engagement variables by providing intuition on engagement patterns and changes to the students' engagement levels over time. Our analysis significantly improves our understanding of the early signs of student drop out.

2 RELATED WORK

Here, we outline related work specifically related to our two contributions: 1) engagement in MOOCs, and 2) predicting grades/dropout/outcomes in online courses. These can be classified into two broad categories: 1) work on classroom and traditional distance education settings, and 2) work on larger settings such as MOOCs.

2.1 Engagement in Classroom Settings

Much of the work before MOOCs concentrate on understanding student engagement using various forms of instructor intervention experiments in classroom settings. Postel et al. [4] analyze the effects of intervention on school dropouts and Tinto et al. [5] examine the reasons behind student attrition in the undergraduate level and discuss possible preventative measures using intervention. Several works perform targeted studies on the effect on intervention on student engagement [6], [7], [8]. Rocca et al. [9] presents an analysis of student engagement in classroom settings, comparing the effects of different methods of teaching on student participation. These studies primarily analyze the effectiveness of various instructor intervention techniques and teaching methodologies on getting students to participate in classroom discussions. Further, these studies primarily refer to participation in classroom discussions as student engagement. Other forms of student engagement such as attending lectures and giving exams are considered integral part of the class. Herrmann [8] analyzes the effect of intervention on passively engaged students to make them engage more actively in the classroom. However, in online settings, the diverse population of the students leads to varied participation levels. This calls for a more nuanced notion of engagement. Drawing analogies from classroom settings and carefully considering student dynamics in online settings, we model three types of student engagement. We refer to participating in discussion forums, which is analogous to participating in classroom discussions as *active engagement*. We refer to following class materials and tests as *passive engagement* and dropping out of the class as *disengagement*. Kuh et al. [10] and Carini et al. [11] study the relationship between student engagement and academic performance for traditional classroom courses; they identify several metrics for user engagement (such as student-faculty interaction, level of academic challenge). Carini et al. [11] demonstrate quantitatively that though most engagement metrics are positively correlated to performance, the relationships in many cases can be weak. Our work borrows ideas from Kuh et al. [10], Carini et al. [11], and from statistical survival models [12] and adapts these to the MOOC setting.

2.2 Engagement in MOOCs

There is growing work studying student engagement in MOOCs [13], [14], [15], [16], [17], [18]. Here, we explain differences of our work from existing work:

- 1) Most existing work only model a single form of engagement and do not differentiate between different forms of engagement such as active and passive [17]. In our work, we model

1. <https://www.coursera.org>

multiple different forms of engagement, active, passive, and absence of engagement as three different variables, thus incorporating the ability to distinguish between these different types of engagement. Also, our engagement variables are continuous-valued, so it is possible for a student to have multiple different types of engagement simultaneously, providing a finer-grained analysis of engagement.

- 2) Our engagement variables are learned via predictive analysis, as opposed to unsupervised models [15], which allow our models to use feedback from student success variables of performance and survival and other features and their combination to guide latent variable values during training.
- 3) We define engagement explicitly according to education theory as discussed by Rocca et al. [9]. The intuitive and interpretable nature of our model that captures dependencies among features and feature-groups and the meaningful nature of our latent engagement variables make our models easy to encode and interpret by domain experts. Existing approaches use machine learning approaches such as logistic regression/factor graphs [13], [16], [19], [20], which lack interpretability on how different features/feature-groups come together to predict student engagement and performance, which our models especially bring forth via first-order logic rules.
- 4) Further, our experimental results in Section 5 demonstrate that our models, especially model with latent engagement variables, can achieve superior prediction performance on courses previously unseen by the model, asserting that the latent engagement variables indeed abstract important behavioral, linguistic, structural, and temporal information that is useful across courses.

2.3 Learning Analytics

There is also a growing body of work in the area of learning analytics. Various works analyze student dropouts in MOOCs [16], [19], [20], [21], [22], [23], [24], [25]. However, all these works only consider final grades as the measure of student success. Due to the presence of a diverse student population in MOOCs, we use a combination of performance and survival for measuring student success. Some works also model student engagement in MOOCs [26], [27], [28], while others focus on discussion forums and post-test performance [29], [30]. These works use students interacting with the online MOOC platform a sign of engagement and analyze the different factors surrounding their online presence such as content in the discussion forums, and quality of the videos. They however do not consider nuanced definitions of engagement that we model in our work. [31] develop models to predict learning outcomes early in online courses. While their approach can predict learning outcomes early, their models function as a black-box classifier, thus providing little insight on how specific features/feature-groups, outcomes, and engagement come together for this prediction. The most significant difference between our approach and existing work on predicting learning outcomes/dropout in MOOCs is that we encode meaningful combinations of several factors that contribute to student engagement and hence their survival in online courses using first-order logic rules, which provide our models with superior interpretability. Further our experimental results show the performance of our models on early prediction and previously unseen courses, which further demonstrates the capabilities of the model in prediction.

Our work will potentially pave the way for constructing better quality MOOCs, which will then result in increase in enrollment and student retention.

2.4 Hinge-loss Markov random fields (HL-MRFs) and Probabilistic Soft Logic

To model the different types of interactions between features and course success, we propose a powerful approach using HL-MRFs. HL-MRFs falls under the class of statistical relational learning models, which combine logic and probability to create richer models. Often in structured domains, first order logic is used to encode intricate dependencies between the different features, latent, and target variables. Statistical relational models use logic to define feature functions in a probabilistic model, to create richer models that are capable of encoding both structural dependencies and uncertainty in the data.

Hinge-loss Markov random fields (HL-MRFs) are a scalable class of continuous, conditional graphical models [32]. Inference of the most probable explanation in HL-MRFs is a convex optimization problem, which makes working with HL-MRFs very efficient in comparison to many relational modeling tools that use discrete representations.

$$P(Y|X) \propto \exp\left(-\sum_{r=1}^M \lambda_r \phi_r(Y, X)\right)$$

$$\phi_r(Y, X) = (\max\{l_r(Y, X), 0\})^{\rho_r}, \quad (1)$$

where $\phi_r(Y, X)$ is a *hinge-loss potential* corresponding to an instantiation of a rule r containing observed features X and target variables Y that we are interested in predicting. The linear function l_r refers to a linear combination of X and Y and an optional exponent $\rho_r \in \{1, 2\}$. λ_r gives the weight of the rule. Each rule is then grounded using actual data creating multiple instantiations of the rule. The weights and potentials are grouped into templates, which are then be used to define HL-MRFs for the MOOC data.

2.4.1 Probabilistic Soft Logic

HL-MRF models can be specified using *Probabilistic Soft Logic (PSL)* [32]. PSL is a framework for collective, probabilistic reasoning in relational domains, which uses syntax based on first-order logic as a templating language for continuous graphical models over random variables representing soft truth values. Like other statistical relational learning methods, PSL uses weighted rules to model the dependencies in a domain. However, one distinguishing aspect is that PSL uses continuous variables to represent truth values, relaxing Boolean truth values to the interval $[0, 1]$. Triangular norms, which are continuous relaxations of logical connectives AND and OR, are used to combine the individual atoms in the first-order clauses. Logical conjunctions of Boolean predicates X and Y ($X \wedge Y$) can be generalized to continuous variables using the hinge function $\max\{X + Y - 1, 0\}$, also known as the Lukasiewicz t-norm. Similarly, disjunctions ($X \vee Y$) are relaxed to $\min\{X + Y, 1\}$, and $\neg X$ to $1 - X$. Using data, we ground out substitutions for these logical terms in the rules. The groundings of a template define hinge-loss potentials that share the same form and the same weight.

An example of a PSL rule is

$$\lambda : P(a) \wedge Q(a, b) \rightarrow R(b),$$

where P , Q , and R are predicates, a and b are *variables*, and λ is the weight associated with the rule. Inference in HL-MRFs is a convex optimization problem, which makes working with PSL very efficient in comparison to relational modeling tools that use discrete representations.

PSL enables us to encode our observed features, latent and target variables as logical predicates and design models by writing rules over these predicates. The expressiveness and flexibility of PSL allows us to easily build different models for MOOC data, and we exploit this by comparing a model that represents multiple forms of latent engagement against a simpler model that directly relates the observable features to student success. To demonstrate this, consider the task of *collectively* predicting student performance, by capturing how students interact with each other in the discussion forums.

Let U_1 and U_2 be two students interacting in the same thread in the discussion forum, posting posts P_1 and P_2 in the discussion forum, respectively. Predicates $\text{POST}(U_1, P_1)$ and $\text{POST}(U_2, P_2)$ denote student U_1 posting P_1 , and U_2 posting P_2 in the discussion forum. The predicate $\text{SAMETHREAD}(P_1, P_2)$ captures if posts P_1 and P_2 are in the same thread. The PSL rule below captures the influence students have on each other when interacting in the forums. Students U_1 and U_2 post in the same threads, hence influence each other to have similar succeeding abilities. This example especially brings out the relational and collective nature of our model, whereby we can reason about users' prediction performance *jointly* based on their interaction with each other.

$$\lambda : \text{POST}(U_1, P_1) \wedge \text{POST}(U_2, P_2) \wedge \text{SAMETHREAD}(P_1, P_2) \wedge \text{SUCCESS}(U_1) \rightarrow \text{SUCCESS}(U_2).$$

The potential $\phi(Y, X) = [\max\{Y_{U_1, P_1}^1 + Y_{U_2, P_2}^1 + Y_{P_1, P_2}^2 + Y_{U_1}^3 - Y_{U_2}^3 - 1, 0\}]^p$ is one minus the truth value of the Boolean formula given above when $Y_{U_1, P_1}^1, Y_{U_2, P_2}^1, Y_{P_1, P_2}^2, Y_{U_1}^3, \text{ and } Y_{U_2}^3 \in [0, 1]$. Since the variables take on values in $[0, 1]$, the potential is a convex relaxation of the implication. An HL-MRF with this potential function assigns higher probability to variable states that satisfy the logical implication above, which can occur to varying degrees in the continuous domain. Given the behavioral data containing all student interactions, PSL constructs the fully ground HL-MRF by grounding out substitutions for different U_1, U_2, P_1 , and P_2 and subsequently generating potential functions for all these substitutions.

2.4.2 Latent Variables in HL-MRFs

HL-MRFs admit various learning algorithms for fully-supervised training data, and are amenable to expectation maximization (EM) for partially-supervised data with latent variables [33]. Latent variables can improve the quality of probabilistic models in many ways. Using latent variables to mediate probabilistic interactions can improve generalization by simplifying models. HL-MRFs' capability in representing continuous latent variables is helpful in expressing more nuanced information when compared to discrete latent variables. Latent variable HL-MRFs are accurate and scalable for three reasons: 1) the continuous variables of HL-MRFs can express complex, latent phenomena, such as mixed group memberships, which add flexibility and modeling power to these models, 2) fast, exact inference for HL-MRFs can identify the most probable assignments to variables quickly, and 3) HL-MRFs can easily express dependencies among latent variables

creating rich, interpretable models. We use this capability to represent student engagement types as a latent variables. We can generate more complex rules connecting the different features and latent variables, which we will demonstrate in Section 3.1.4. The HL-MRF model uses these rules to encode domain knowledge about dependencies among the predicates. The continuous value representation further helps in understanding the confidence of predictions. In Section 3.1, we detail the various features we collect from the data.

3 STUDENT SUCCESS PREDICTION MODELS

As students interact on a MOOC, detailed records are generated, including page and video views, forum visits, forum interactions such as voting, posting messages and replies, and graded elements such as quizzes and assignments. In this section, we develop our models for predicting student success in MOOCs. Our models connect performance indicators to complex behavioral, linguistic, temporal, and structural features derived from the raw student interactions. Our first model, referred as the DIRECT model, directly encodes the dependence between student interactions and student success in MOOCs. We then extend the DIRECT model by adding latent variables modeling three types of student engagement: 1) active engagement, 2) passive engagement, and 3) disengagement. We refer to this model as the LATENT model. In the LATENT model, we capture dependencies among student interactions, their different types of engagement, and success measures.

We evaluate the models by employing them to predict student success in MOOCs. We consider two course success indicators in MOOCs: 1) *performance*: whether the student earns a certificate in the course, and 2) *survival*: whether the student follows the course till the end.

3.1 Modeling MOOC Student Activity

MOOC students interact with three main resources on the MOOC website: video lectures, quizzes, and discussion forums. Students can watch lectures multiple times and respond to on-demand quizzes during the lectures. Students can interact by asking and responding to questions in the forums. There are typically multiple forums organized by topics, each consisting of multiple threads, and each thread consisting of multiple posts. Students can respond, vote (up or down) on existing posts and subscribe for updates to forums threads. Each student is given a reputation score based on the votes on posts created by the student. These activities are depicted in Figure 1. Though our datasets are all from Coursera, the core activities captured in Figure 1 are present in all other MOOCs offered by other popular companies such as EdX and Udacity; they also have video lectures, quizzes and discussion forum posts and ability to view, follow, reply to, and upvote/downvote discussion forum posts, making our features extensible across platforms.

We quantify these activities by defining a set of PSL predicates over the raw student data, and capture more complex behaviors by combining these predicates into expressive rules, used as features in our predictive models. We categorize these predicates as either behavioral, linguistic, structural, or temporal, and describe them in the following sections.

3.1.1 Behavioral Features

Behavioral features are derived from various activities that students engage in while interacting on the MOOC website. These

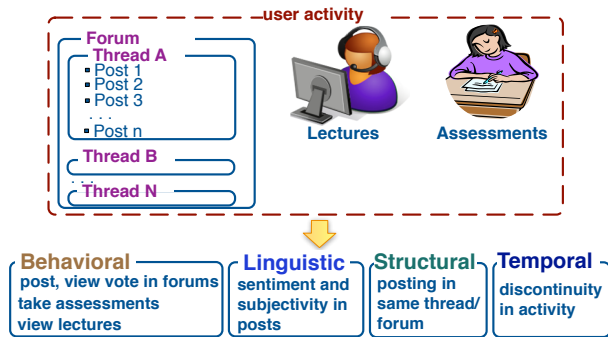


Fig. 1: Structure of MOOC student activity.

features measure the different levels of activity of MOOC participants on the site. We consider three types of student interactions on the discussion forums: posting in the forums, voting on forum posts, and viewing forum posts. We consider two types of behavioral features: aggregate and non-aggregate. Aggregate features are predicates comparing students' activity level to the median value of that activity considering all students. With the median value of student activity corresponding to a value of 0.5 for the predicate, all other values are scaled appropriately to have a value in $(0,1)$. The predicates $\text{POST-ACTIVITY}(\text{USER})$, $\text{VOTE-ACTIVITY}(\text{USER})$ and $\text{VIEW-ACTIVITY}(\text{USER})$ represent aggregate features capturing student activity in the forums. Non-aggregate features directly quantify student's behavior. The predicates $\text{POSTS}(\text{USER}, \text{POST})$ and $\text{VOTES}(\text{USER}, \text{POST})$ capture an instance-level log of users posting and voting on the discussion forums. The predicates POSTS and VOTES are true if the USER posts or votes on POST . Predicate $\text{UPVOTE}(\text{POST})$ is true if the post has positive votes and false otherwise, and predicate $\text{DOWN-VOTE}(\text{POST})$ is true if a post has been down-voted. In addition to that, we also measure the reputation of student in the forum taking into account, the total number of upvotes/downvotes gained by the student across all the posts. We refer to this aggregate feature as $\text{REPUTATION}(\text{USER})$ in our model. The student who gathers the most upvotes gets a score of 1.0 and the student who gathers the most downvotes gets a score of 0.0 and all other students get a score in $(0, 1)$.

The second class of behavioral features capture students' interaction with lectures and quizzes on the MOOC website. We measure the percentage of lectures and accompanying quizzes that were submitted by the student in the course. The features $\text{LECTURE-VIEWED}(\text{USER})$ captures the fraction of lectures submitted by the student in the course. The feature $\text{LECTURE-VIEWED-ONTIME}(\text{USER})$ captures the fraction of lectures submitted by the student within the due date. Similarly, for quizzes we derive QUIZ-SUBMITTED and $\text{QUIZ-SUBMITTED-ONTIME}(\text{USER})$. These predicates are continuous valued in $[0, 1]$.

3.1.2 Forum Content and Interaction Features

MOOC forums are rich with relevant information, indicative of the students' attitudes toward the course and its materials as well as the social interactions between students. We capture this information using two types of features, *linguistic* features capturing the sentiment of the post content, and *structural* features

capturing the forum structure, organized topically into threads and forums types.

3.1.2.1 Linguistic Features: The attitudes expressed by students on the forums can be captured by estimating sentiment polarity (positive or negative) and identifying subjective posts. Since MOOC forums contain thousands of posts, we use an automated tool, *OpinionFinder* [34] to avoid manual annotation. The tool segments the forums posts into sentences, and assigns subjectivity and polarity tags for each sentence. Based on its predictions, we define two predicates, $\text{POLARITY}(\text{POST})$ and $\text{SUBJECTIVE}(\text{POST})$. Both predicates are calculated by normalizing the number of subjective/objective tags and positive/negative polarity tags marked by *OpinionFinder*. The normalization keeps these values in the $[0, 1]$ interval, where values close to 0.0 indicate that the post has negative polarity and values close to 1.0 indicate that the post has positive polarity.

Tables 1 show some examples of posts having negative polarity and positive polarity scores. Most negative sentiment posts in MOOC forums are on logistic issues as evidenced in Table 1. Posts that get a value around 0.5 are either neutral posts or posts with both positive and negative sentiment words (Table 2). Positive sentiment posts mostly are either feedback posts or posts that thank the instructor or other students when they respond to their queries. In our models, we especially focus on positive and negative polarity posts as indicated by $\text{POLARITY}(\text{POST})$ and $\neg\text{POLARITY}(\text{POST})$.

3.1.2.2 Structural Features: Forums are structured entities, organized by high-level topics (at the forum level) and specific topics (thread level). Including these structural relationships allows our model to identify structural relations between forum posts and connect them with students participating in the forum discussions. The predicates representing forum structure are $\text{SAME-THREAD}(\text{POST}_1, \text{POST}_2)$ and $\text{SAME-FORUM}(\text{THREAD}_1, \text{THREAD}_2)$, which are true for posts in the same thread and threads in the same forum, respectively. These predicates capture forum interaction among students and propagate *performance*, *survival* and *engagement* values among them. Table 3 gives posts from some example threads. We observe that posts in the same thread often contain posts on topics that have certain amount of connectivity as considered by [35]. Even if this is not the case, the students posting on the same threads, may have a certain amount of overlap in interests. In our rules, we model this interaction and how it influences their respective survival capabilities using the SAME-THREAD and SAME-FORUM predicates. These rules also help us use behavioral and interaction features from students to have strong signals to infer performance, survival, and engagement values for students who have less behavioral information. For example, in Table 3, we find that post 1 and 2 are both reporting the same issue. Looking closely at the posts, both the students seem to be interested in completing the assignment and are likely to have similar performance and survival. So it is possible to improve prediction accuracy for the students based on the features and prediction of the other student.

3.1.3 Temporal Features

Student activity levels change over the span of the course. Students are often active at early stages and lose interest as the course progresses. To include signals of how student activity changes over time, we introduce a set of temporal features. We divide the course into three time periods: *start*, *mid*, and *end*. The time period splits are constructed by dividing the course by duration into three

polarity	example post
polarity = 0.25	JSTOR allowed 3 items (texts/writings) on my 'shelf' for 14 days. But, I read the items and wish to return them, but cannot, until 14 days has expired. It is difficult then, to do the extra readings in the "Exploring Further" section of Week 1 reading list in a timely manner. Does anyone have any ideas for surmounting this issue?
polarity = 0.0	There are some mistakes on quiz 2. Questions 3, 5, and 15 mark you wrong for answers that are correct.
polarity = 0.9	Kudos to the Professor for a great course!

TABLE 1: Negative sentiment posts

polarity = 0.45	This course is very interesting. I initially had some trouble, but managed to do well.
polarity = 0.4	I am sort of disappointed that my final grade did not turn out to be that good. But I enjoyed the course and look forward to the next course in the sequence.

TABLE 2: Posts having both negative and positive sentiment

polarity = 0.0	I was just looking at the topics for the second essay assignments. The thing is I don't see what the question choices are. I have the option of Weeks and I have no idea what that even means. Can someone help me out here and tell me what the questions for the second essay assignment are I think my computer isn't allowing me to see the whole assignment! Someone please help me out and let me know that the options are.
polarity = 0.25	I'd appreciate someone looks into this at the earliest. I am having the same problem with the essay assignments. Thanks..
polarity = 0.78	Hopefully the essay assignments now open for you. Thanks for reporting this.

TABLE 3: Example posts in a thread

equal chunks. The temporal features LAST-QUIZ, LAST-LECTURE, LAST-POST, LAST-VIEW and LAST-VOTE indicate the time-period in which each last interaction of the user occurred. These features measure to what lengths the user participated in different aspects of the course.

3.1.4 Constructing Complex Rules

We use the features above to construct PSL rules using logical connectives, as demonstrated in Table 4. We construct meaningful combinations of predicates to model student engagement and student success. Our rules combine features across the different feature categories, discrete and continuous feature values, and observed, latent, and target variables to capture intricate dependencies in the data. For example, the first rule in Table 4 combines the posting activity of user U relative to other students in the class (POST-ACTIVITY) with reputation of the user in the forums to infer student success. This rule captures that students posting high-quality posts (given by reputation) show greater signs of succeeding in the class. This is helpful in discerning between students who post a lot and students who post few highly upvoted posts. Similarly, the third rule combines posting in forums and the polarity of forum posts to capture that students posting positive sentiment posts are more likely to engage and succeed in the course. The PSL models associate these rules with student success, either directly or indirectly using latent variables. We explain this process in Section 4.

3.2 Student Engagement in MOOCs

Student engagement cannot be directly measured from the data. The interpretable nature of our models (i.e., encoded in first order logic) makes it possible to abstract definitions of engagement in latent engagement variables using combinations of observed features and student success target variables. We therefore treat student engagement as latent variables and associate various observed features to one or more forms of engagement. Drawing analogies from classroom settings and adapting them to the online

<ul style="list-style-type: none"> • Behavioral Features
$POST-ACTIVITY(U) \wedge REPUTATION(U) \rightarrow SUCCESS(U)$
$LECTURE-VIEWED(U) \wedge LECTURE-VIEWED-ONTIME(U) \rightarrow SUCCESS(U)$
<ul style="list-style-type: none"> • Forum Content Features
$POSTS(U, P) \wedge POLARITY(P) \rightarrow SUCCESS(U)$
$POSTS(U, P) \wedge \neg POLARITY(P) \rightarrow \neg SUCCESS(U)$
<ul style="list-style-type: none"> • Forum Interaction Feature
$POSTS(U_1, P_1) \wedge POSTS(U_2, P_2) \wedge SAME-THREAD(P_1, P_2) \rightarrow SUCCESS(U)$
<ul style="list-style-type: none"> • Temporal Features
$LAST-QUIZ(U, T_1) \wedge LAST-LECTURE(U, T_1) \wedge LAST-POST(U, T_1) \rightarrow SUCCESS(U)$

TABLE 4: Constructing complex rules in PSL

settings, we model three types of student engagement. These three types of engagement are denoted by three engagement variables, ACTIVE-ENGAGEMENT, PASSIVE-ENGAGEMENT and DISENGAGEMENT. ACTIVE-ENGAGEMENT represents students actively engaged in the course by participating in the forums, PASSIVE-ENGAGEMENT represents students following the class materials but not making an active presence in the forums, and DISENGAGEMENT represents students discontinuing from engaging with the course both actively or passively. We associate different features representing MOOC attributes relevant for each engagement type. Our engagement scores for each student across the three types of engagement are normalized to sum to 1.

- **Active Engagement** Actively participating in course-related discussions by posting in the forums are signs of active engagement.
- **Passive Engagement** Passively following course material by viewing lectures, viewing/voting/subscribing to posts on discussion forums, and giving quizzes are signs of passive engagement.
- **Disengagement** Temporal features, indicating the last point of user's activity, capture signs of disengagement.

4 PSL MODELS FOR STUDENT SUCCESS PREDICTION

We construct two different PSL models for predicting student success in a MOOC setting—first, a model (denoted **DIRECT**) that directly infers student success from observable features, and second, a latent variable model (**LATENT**) that infers student engagement as a hidden variable to predict student success. By building both models, we are able to evaluate the contribution of the abstraction created by formulating engagement patterns as latent variables.

4.1 PSL-DIRECT

In **PSL-DIRECT** model, we model student success by using the observable behavioral features exhibited by the student, linguistic features corresponding to the content of posts, structural features derived from forum interactions, and temporal features capturing discontinuity in activity. Meaningful combinations of one or more observable behavioral, linguistic, temporal, and structural features are constructed as described in Section 3.1 and they are used to predict student **SUCCESS**. Table 5 contains the rules used in the **DIRECT** model. **U** and **P** in tables 5, 6, and 7 refer to **USER** and **POST** respectively. The **DIRECT** model rules allow observable features to directly imply student success. For ease of understanding, we categorize the rules into four groups based on the features present in them. The first group of rules presents the different combinations of student interactions with the three course elements: discussion forums, lectures, and quizzes, to predict student success indicated by **SUCCESS**. Note that we capture combinations of features to infer student success. For example, the fourth rule in the first group combines posting activity, viewing activity, and voting activity to infer student success. Similarly, we combine viewing lectures (**VIEW-LECTURE**) and if they were viewed before the due date (**ONTIME**) to infer success. We use a similar combination for quizzes as well combining taking quizzes (**SUBMITTED-QUIZ**) and the taking them before the due date (**ONTIME-QUIZ**) to infer student success. The second group of rules combine the behavioral features with the linguistic features to predict student success. Here, we combine posting on the forums, which is a behavioral feature with the linguistic features such as polarity of the post, to infer student success. The third set of rules capture the structural interactions of students with other fellow students in the forums and how that impacts each other's course succeeding capabilities. The last set of rules capture the interaction between behavioral and temporal features.

4.2 PSL-LATENT

In the **LATENT** model, we enhance reasoning in the **DIRECT** model by including latent variables semantically based on concepts of student engagement as outlined in Section 3.2. We introduce three latent variables **ACTIVE-ENGAGEMENT**, **PASSIVE-ENGAGEMENT**, and **DISENGAGEMENT** to capture the three different types of student engagement. We present the **LATENT** model in two parts in Tables 6 and 7. In Table 6, we present rules connecting observable features to different forms of engagement. It is important to note that both our models have been provided the same set of features. Also, note that the rules in the **LATENT** model are identical to the rules in the **DIRECT** model presented in Table 5, except that in the **LATENT** model they are changed to imply the latent engagement variables instead of student success.

In this model, some of the observable features (e.g. **POST-ACTIVITY**, **VOTE-ACTIVITY**, **VIEW-ACTIVITY**) are used to classify students into one or more forms of engagement or disengagement. For example, in Table 6, conjunction of **POST-ACTIVITY** and **REPUTATION** implies **ACTIVE-ENGAGEMENT**; conjunction of **VOTE-ACTIVITY** and **REPUTATION** implies **PASSIVE-ENGAGEMENT**. Rules that combine observed features that are indicative of more than one form of engagement, such as **POST-ACTIVITY** and **VOTEACTIVITY**, are left unchanged from the **DIRECT** model to directly imply **SUCCESS**. We then connect the latent engagement variables to student success using the rules in Table 7. For example, **ACTIVE-ENGAGEMENT** and **PASSIVE-ENGAGEMENT** implies **SUCCESS**. We consider various combinations of engagement and their relationship to **SUCCESS**. For example, exhibiting both passive and active forms of engagement implies **SUCCESS**. Also, exhibiting only one form of engagement, either active or passive, implies **SUCCESS**. In Section 5, we present results from training and testing our models on the two success measures. The resulting model with latent engagement suggests which forms of engagement are good indicators of student success. We demonstrate that the **LATENT** model not only produces better predictive performance, but also provides more insight into MOOC user behavior when compared to the **DIRECT** model.

4.3 Weight Learning

We train the weights for both the models using **SUCCESS** as the target variable. The weighted combinations of different engagement types encodes variations in student engagement types and their relationship to student success. The weights of the rules in the **PSL-DIRECT** model are learned by maximum likelihood estimation. This is accomplished by finding the parameter values (weight values) that will maximize the likelihood of the data given the parameters. In the **PSL-LATENT** model, due to the presence of latent variables, the rule weights are learned by performing expectation maximization (**EM**), which iterates alternatively between estimating the values of the latent variables and weight values till a local optimum solution is achieved. This is carried out by first estimating the expected value of the latent engagement variables in the current setting of the weights. Then, using the estimated expected values of latent variables and the ground truth values of target outcome variables, the new weights are estimated by finding the values of the parameters that will maximize the likelihood of the data given the parameter values.

5 EMPIRICAL EVALUATION

Here, we present our detailed experimental evaluation of our models. We conduct extensive experiments to answer the following questions.

- 1) How effective are our models at predicting student success: performance and survival in online courses?
- 2) How effective are our models at predicting student survival considering student interactions only from early part of the course?
- 3) How effective are our models at predicting student survival on previously unseen courses and how reliably can they predict student survival on unseen courses by considering student interactions from only the early part of the course?
- 4) How useful are our different classes of features in predicting student success, across different time periods in the course?
- 5) How useful are the values learned by the latent engagement variables?

PSL-DIRECT RULES

Rules combining behavioral features

$\text{POST-ACTIVITY}(U) \wedge \text{REPUTATION}(U) \rightarrow \text{SUCCESS}(U)$
 $\text{VOTE-ACTIVITY}(U) \wedge \text{REPUTATION}(U) \rightarrow \text{SUCCESS}(U)$
 $\text{VIEW-ACTIVITY}(U) \wedge \text{REPUTATION}(U) \rightarrow \text{SUCCESS}(U)$
 $\text{POST-ACTIVITY}(U) \wedge \text{VIEW-ACTIVITY}(U) \wedge \text{VOTE-ACTIVITY}(U) \rightarrow \text{SUCCESS}(U)$
 $\neg \text{POST-ACTIVITY}(U) \rightarrow \neg \text{SUCCESS}(U)$
 $\neg \text{VOTE-ACTIVITY}(U) \rightarrow \neg \text{SUCCESS}(U)$
 $\neg \text{VIEW-ACTIVITY}(U) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{POST-ACTIVITY}(U) \wedge \neg \text{REPUTATION}(U) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{POSTS}(U, P) \wedge \text{REPUTATION}(U) \rightarrow \text{SUCCESS}(U)$
 $\text{VIEWED-LECTURE}(U) \rightarrow \text{SUCCESS}(U)$
 $\neg \text{VIEWED-LECTURE}(U) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{VIEWED-LECTURE}(U) \wedge \text{ONTIME}(U) \rightarrow \text{SUCCESS}(U)$
 $\text{VIEWED-LECTURE}(U) \wedge \neg \text{ONTIME}(U) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{SUBMITTED-QUIZ}(U) \rightarrow \text{SUCCESS}(U)$
 $\neg \text{SUBMITTED-QUIZ}(U) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{SUBMITTED-QUIZ}(U) \wedge \text{ONTIME-QUIZ}(U) \rightarrow \text{SUCCESS}(U)$
 $\text{SUBMITTED-QUIZ}(U) \wedge \neg \text{ONTIME-QUIZ}(U) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{SUBMITTED-QUIZ}(U) \wedge \text{SUBMITTED-QUIZ}(U) \rightarrow \text{SUCCESS}(U)$

Rules combining behavioral and linguistic features

$\text{POSTS}(U, P) \wedge \text{POLARITY}(P) \rightarrow \text{SUCCESS}(U)$
 $\text{POSTS}(U, P) \wedge \neg \text{POLARITY}(P) \rightarrow \neg \text{SUCCESS}(U)$

Rules combining behavioral and structural features

$\text{POSTS}(U_1, P_1) \wedge \text{POSTS}(U_2, P_2) \wedge \text{SUCCESS}(U_1) \wedge \text{SAME-THREAD}(P_1, P_2) \rightarrow \text{SUCCESS}(U_2)$
 $\text{POSTS}(U_1, P_1) \wedge \text{POSTS}(U_2, P_2) \wedge \text{SUCCESS}(U_1) \wedge \text{SAME-FORUM}(P_1, P_2) \rightarrow \text{SUCCESS}(U_2)$

Rules combining behavioral and temporal features

$\text{LAST-POST}(U, \textit{start}) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{LAST-LECTURE}(U, \textit{start}) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{LAST-QUIZ}(U, \textit{start}) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{LAST-POST}(U, \textit{mid}) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{LAST-LECTURE}(U, \textit{mid}) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{LAST-QUIZ}(U, \textit{mid}) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{LAST-POST}(U, \textit{end}) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{LAST-LECTURE}(U, \textit{end}) \rightarrow \text{SUCCESS}(U)$
 $\text{LAST-LECTURE}(U, \textit{end}) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{LAST-QUIZ}(U, \textit{end}) \rightarrow \text{SUCCESS}(U)$
 $\text{LAST-QUIZ}(U, \textit{end}) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{LAST-QUIZ}(U, \textit{end}) \wedge \text{LAST-LECTURE}(U, \textit{end}) \wedge \text{LAST-POST}(U, \textit{end}) \rightarrow \text{SUCCESS}(U)$
 $\text{LAST-QUIZ}(U, \textit{end}) \wedge \text{LAST-LECTURE}(U, \textit{end}) \wedge \text{LAST-POST}(U, \textit{end}) \rightarrow \neg \text{SUCCESS}(U)$

TABLE 5: Rules from the PSL-DIRECT model

5.1 Datasets and Experimental Setup

We evaluate our models on seven Coursera MOOCs at University of Maryland: *Surviving Disruptive Technologies*, *Women and the Civil Rights Movement*, two iterations of *Gene and the Human Condition*, and three iterations of *Developing Innovative Ideas for New Companies*. These courses cover a broad spectrum of topics spanning across humanities, business, and sciences. We refer to these courses as DISR, WOMEN, GENE-1, GENE-2, INNO-1, INNO-2 and INNO-3, respectively. DISR is 4 weeks, WOMEN is 5 weeks, GENE is 8 weeks, and INNO is 4 weeks in duration. Our data consists of anonymized student records, grades, and online behavior recorded during each course duration.

Figure 2 shows the number of participants in different course-related activities. Of the total number of students registered, around 5% of the students in DISR-TECH and WOMEN, 14% in GENE-1, 21% in GENE-2, 7% in INNO-1, 15% in INNO-2, and 5% in INNO-3 complete the course. In all the courses, the most prominent activity exhibited by students while on the site is viewing lectures. Hence, we rank students based on number of lectures viewed, as a baseline (denoted LECTURE-RANK in our tables) for comparison. The other prevalent activities include submitting

quizzes and viewing forum content. Observing the statistics, DISR and WOMEN have a higher percentage of total registered students participating in forums compared to GENE and INNO courses. We also run various classical machine learning models (SVM, Logistic Regression, Multi-layer Perceptron, Linear Regression, Decision Trees) using all the features included in our model except the features that these models are not capable of representing (structural features) and compare against against the best performing one (indicated as classical ML model in Tables 8 and 9). These models use all the features except structural features that capture specific structural relationships among different users/posts that are unique to statistical relational models such as HL-MRFs.

We evaluate the model on the following metrics: area under the precision-recall curve for positive and negative labels and area under the ROC curve. We use ten-fold cross-validation, leaving out 10% of the data for testing and revealing the rest for training the model weights. Statistically significant differences, evaluated using a paired t-test with a rejection threshold of 0.01, are typed in bold.

PSL-LATENT RULES (PART 1)

Rules combining behavioral features

$\text{POST-ACTIVITY}(U) \wedge \text{REPUTATION}(U) \rightarrow \text{ACTIVE-ENGAGEMENT}(U)$
 $\text{VOTE-ACTIVITY}(U) \wedge \text{REPUTATION}(U) \rightarrow \text{PASSIVE-ENGAGEMENT}(U)$
 $\text{VIEW-ACTIVITY}(U) \wedge \text{REPUTATION}(U) \rightarrow \text{PASSIVE-ENGAGEMENT}(U)$
 $\text{POST-ACTIVITY}(U) \wedge \text{VIEW-ACTIVITY}(U) \wedge \text{VOTE-ACTIVITY}(U) \rightarrow \text{SUCCESS}(U)$
 $\text{REPUTATION} \rightarrow \text{ACTIVE-ENGAGEMENT}(U)$
 $\neg \text{POST-ACTIVITY}(U) \rightarrow \neg \text{ACTIVE-ENGAGEMENT}(U)$
 $\neg \text{VOTE-ACTIVITY}(U) \rightarrow \neg \text{PASSIVE-ENGAGEMENT}(U)$
 $\neg \text{VIEW-ACTIVITY}(U) \rightarrow \neg \text{PASSIVE-ENGAGEMENT}(U)$
 $\text{POST-ACTIVITY}(U) \wedge \neg \text{REPUTATION}(U) \rightarrow \neg \text{ACTIVE-ENGAGEMENT}(U)$
 $\text{POSTS}(U, P) \wedge \text{REPUTATION}(U) \rightarrow \text{ACTIVE-ENGAGEMENT}(U)$
 $\text{VIEWED-LECTURE}(U) \rightarrow \text{PASSIVE-ENGAGEMENT}(U)$
 $\neg \text{VIEWED-LECTURE}(U) \rightarrow \neg \text{PASSIVE-ENGAGEMENT}(U)$
 $\text{VIEWED-LECTURE}(U) \wedge \text{ONTIME}(U) \rightarrow \text{PASSIVE-ENGAGEMENT}(U)$
 $\text{VIEWED-LECTURE}(U) \wedge \neg \text{ONTIME}(U) \rightarrow \neg \text{PASSIVE-ENGAGEMENT}(U)$
 $\text{VIEWED-LECTURE}(U) \wedge \text{POST-ACTIVITY}(U) \rightarrow \text{PASSIVE-ENGAGEMENT}(U)$
 $\text{SUBMITTED-QUIZ}(U) \rightarrow \text{PASSIVE-ENGAGEMENT}(U)$
 $\text{SUBMITTED-QUIZ}(U) \rightarrow \neg \text{PASSIVE-ENGAGEMENT}(U)$
 $\text{SUBMITTED-QUIZ}(U) \wedge \text{ONTIME-QUIZ}(U) \rightarrow \text{PASSIVE-ENGAGEMENT}(U)$

Rules combining behavioral and linguistic features

$\text{POSTS}(U, P) \wedge \text{POLARITY}(P) \rightarrow \text{ACTIVE-ENGAGEMENT}(U)$
 $\text{POSTS}(U, P) \wedge \neg \text{POLARITY}(P) \rightarrow \neg \text{ACTIVE-ENGAGEMENT}(U)$

Rules combining behavioral and structural features

$\text{POSTS}(U_1, P_1) \wedge \text{POSTS}(U_2, P_2) \wedge \text{ACTIVE-ENGAGEMENT}(U_1) \wedge \text{SAME-THREAD}(P_1, P_2) \rightarrow \text{ACTIVE-ENGAGEMENT}(U_2)$
 $\text{POSTS}(U_1, P_1) \wedge \text{POSTS}(U_2, P_2) \wedge \text{ACTIVE-ENGAGEMENT}(U_1) \wedge \text{SAME-FORUM}(P_1, P_2) \rightarrow \text{ACTIVE-ENGAGEMENT}(U_2)$

Rules combining behavioral and temporal features

$\text{LAST-POST}(U, \textit{start}) \rightarrow \text{DISENGAGEMENT}(U)$
 $\text{LAST-LECTURE}(U, \textit{start}) \rightarrow \text{DISENGAGEMENT}(U)$
 $\text{LAST-QUIZ}(U, \textit{start}) \rightarrow \text{DISENGAGEMENT}(U)$
 $\text{LAST-POST}(U, \textit{mid}) \rightarrow \text{DISENGAGEMENT}(U)$
 $\text{LAST-LECTURE}(U, \textit{mid}) \rightarrow \text{DISENGAGEMENT}(U)$
 $\text{LAST-QUIZ}(U, \textit{mid}) \rightarrow \text{DISENGAGEMENT}(U)$
 $\text{LAST-POST}(U, \textit{end}) \rightarrow \text{DISENGAGEMENT}(U)$
 $\text{LAST-POST}(U, \textit{end}) \rightarrow \text{ACTIVE-ENGAGEMENT}(U)$
 $\text{LAST-LECTURE}(U, \textit{end}) \rightarrow \text{DISENGAGEMENT}(U)$
 $\text{LAST-LECTURE}(U, \textit{end}) \rightarrow \text{PASSIVE-ENGAGEMENT}(U)$
 $\text{LAST-QUIZ}(U, \textit{end}) \rightarrow \text{DISENGAGEMENT}(U)$
 $\text{LAST-QUIZ}(U, \textit{end}) \rightarrow \text{PASSIVE-ENGAGEMENT}(U)$
 $\text{LAST-QUIZ}(U, \textit{end}) \wedge \text{LAST-LECTURE}(U, \textit{end}) \wedge \text{LAST-POST}(U, \textit{end}) \rightarrow \text{SUCCESS}(U)$
 $\text{LAST-QUIZ}(U, \textit{end}) \wedge \text{LAST-LECTURE}(U, \textit{end}) \wedge \text{LAST-POST}(U, \textit{end}) \rightarrow \neg \text{SUCCESS}(U)$

TABLE 6: Rules from the PSL-LATENT model capturing dependencies between observed features and latent engagement variables

PSL-LATENT RULES (PART 2)

Rules combining latent engagement variables

$\text{PASSIVE-ENGAGEMENT}(U) \rightarrow \text{SUCCESS}(U)$
 $\neg \text{PASSIVE-ENGAGEMENT}(U) \rightarrow \neg \text{SUCCESS}(U)$
 $\text{ACTIVE-ENGAGEMENT} \rightarrow \text{SUCCESS}(U)$
 $\neg \text{ACTIVE-ENGAGEMENT} \rightarrow \neg \text{SUCCESS}(U)$
 $\text{PASSIVE-ENGAGEMENT}(U) \wedge \text{ACTIVE-ENGAGEMENT} \rightarrow \text{SUCCESS}(U)$
 $\text{PASSIVE-ENGAGEMENT}(U) \wedge \neg \text{ACTIVE-ENGAGEMENT} \rightarrow \text{SUCCESS}(U)$
 $\text{PASSIVE-ENGAGEMENT}(U) \wedge \neg \text{ACTIVE-ENGAGEMENT} \rightarrow \neg \text{SUCCESS}(U)$
 $\neg \text{PASSIVE-ENGAGEMENT}(U) \wedge \text{ACTIVE-ENGAGEMENT} \rightarrow \text{SUCCESS}(U)$
 $\neg \text{PASSIVE-ENGAGEMENT}(U) \wedge \text{ACTIVE-ENGAGEMENT} \rightarrow \neg \text{SUCCESS}(U)$
 $\neg \text{PASSIVE-ENGAGEMENT}(U) \wedge \neg \text{ACTIVE-ENGAGEMENT} \rightarrow \neg \text{SUCCESS}(U)$
 $\text{DISENGAGEMENT} \rightarrow \neg \text{SUCCESS}(U)$

TABLE 7: Rules from the PSL-LATENT model capturing dependencies between latent engagement variables and student success




Fig. 2: Comparison of number of students participating in course-related activities in seven courses.

5.2 Student Performance Analysis

We conduct experiments to assess how effective our models are in predicting student performance, as measured both by their official grade and whether they complete the course requirements. We also look at the key factors influencing student performance in the online setting as determined by our model. We filter the dataset to include only students that participated in at least one of the possible course related activities. For these students, we label the ones who earn a certificate from the course as positive instances ($PERFORMANCE = 1.0$) and students that did not as negative instances ($PERFORMANCE = 0.0$). In our datasets, we observe that the percentage of students with performance = 1.0 is around 40 – 50% of the filtered set of students. These labels are used as ground truth to train and test the models. Our experimental results are summarized in Tables 8, and show performance values for the DIRECT and LATENT PSL models compared to the LECTURE-RANK and CLASSICAL ML MODEL baseline. We observe that the LATENT PSL model performs better at predicting students performance, outperforming both the DIRECT, LECTURE-RANK, and CLASSICAL ML MODEL models.

To better understand which behavioral factors provide more predictive information, we examine the weights our models learned at training time. The rules involving viewing lectures and viewing forum posts have highest weights in the DIRECT learned model, indicating the importance of these features in predicting performance. The other prominent features which get high weights in the learned model are posting in forums, and reputation of student in the forums. In the LATENT model, rules corresponding to passive engagement have highest weights in the learned model for predicting performance. This emphasizes the importance of passive forms of engagement in online settings. This is followed by rules corresponding to active engagement, indicating that active forms of engagement are also predictive of student success in online courses, but fall second to passive forms of engagement. Rules corresponding to disengagement gain high weights for predicting student drop out.

5.3 Student Survival Analysis

Our experiments in the student survival models are aimed at measuring student survival by understanding factors influencing students’ survival in the course, engagement types and changes in engagement, and the effectiveness of prediction at different time periods of the course. For survival analysis, we consider all registered students in the course. We observe that the percentage of survived students is around 5 – 10% in the total number of students. Note that while we filter students based on their activity

COURSE	MODEL	AUC-PR Pos.	AUC-PR Neg.	AUC-ROC
DISR	LECTURE-RANK	0.630	0.421	0.512
	CLASSICAL ML MODEL	0.397	0.623	0.505
	DIRECT	0.739	0.546	0.667
	LATENT	0.749	0.575	0.692
WOMEN	LECTURE-RANK	0.263	0.761	0.503
	CLASSICAL ML MODEL	0.260	0.769	0.521
	DIRECT	0.557	0.881	0.767
	LATENT	0.732	0.959	0.909
GENE-1	LECTURE-RANK	0.503	0.482	0.476
	CLASSICAL ML MODEL	0.476	0.528	0.499
	DIRECT	0.814	0.755	0.817
	LATENT	0.943	0.879	0.931
GENE-2	LECTURE-RANK	0.466	0.522	0.482
	CLASSICAL ML MODEL	0.491	0.528	0.512
	DIRECT	0.806	0.783	0.831
	LATENT	0.923	0.941	0.932
INNO-1	LECTURE-RANK	0.376	0.651	0.507
	CLASSICAL ML MODEL	0.380	0.621	0.501
	DIRECT	0.714	0.858	0.815
	LATENT	0.850	0.920	0.899
INNO-2	LECTURE-RANK	0.536	0.984	0.938
	CLASSICAL ML MODEL	0.545	0.530	0.537
	DIRECT	0.785	0.790	0.811
	LATENT	0.892	0.876	0.881
INNO-3	LECTURE-RANK	0.239	0.813	0.543
	CLASSICAL ML MODEL	0.240	0.799	0.533
	DIRECT	0.586	0.930	0.835
	LATENT	0.833	0.983	0.945

TABLE 8: Performance of LECTURE-RANK, DIRECT and LATENT models in predicting student performance

for predicting performance, here we apply no filtering and consider all students enrolled in the course. By not filtering the students based on their activity enables our models to be used directly off-the-shelf for predicting survival without the need for any pre-processing. As can be observed from Figure 2, a high proportion of students drop out from MOOCs, leading to a huge class imbalance in the data. By using a combination of filtering (for predicting performance) and no filtering (for predicting survival), we demonstrate the utility of our models in two settings: i) when there is little or no class imbalance, and ii) when class imbalance is present. Due to the huge class imbalance in the data, models that can identify students who will survive the course are more valuable in this setting. The LECTURE-RANK and CLASSICAL ML MODEL baselines can predict dropouts reasonably well, but its comparatively low precision and recall for positive survival (AUC-PR pos.), with CLASSICAL ML MODEL sometimes performing worse than LECTURE-RANK, indicates that using these models are suboptimal for predicting survival. We consider all student activity during the entire course to predict whether each student takes the final quiz. The scores for our DIRECT and LATENT survival models, CLASSICAL ML MODEL, and LECTURE-RANK baselines are listed in Table 9. The strength of our models comes from combining behavioral, linguistic, temporal, and structural features for predicting student survival. Our models DIRECT and LATENT significantly improve on the baselines, and the LATENT model outperforms the DIRECT model.

5.4 Early Survival Prediction

Predicting student survival can provide instructors with a powerful tool if these predictions can be made reliably before the students disengage and drop out. We simulate this scenario by training our model over data collected early in the course. We divide the course

COURSE	MODEL	AUC-PR Pos.	AUC-PR Neg.	AUC-ROC
DISR	LECTURE-RANK	0.333	0.998	0.957
	CLASSICAL ML MODEL	0.343	0.998	0.957
	DIRECT	0.393	0.997	0.936
	LATENT	0.546	0.998	0.969
WOMEN	LECTURE-RANK	0.508	0.995	0.946
	CLASSICAL ML MODEL	0.049	0.951	0.500
	DIRECT	0.565	0.995	0.940
	LATENT	0.816	0.998	0.983
GENE-1	LECTURE-RANK	0.688	0.984	0.938
	CLASSICAL ML MODEL	0.139	0.861	0.500
	DIRECT	0.793	0.997	0.976
	LATENT	0.818	0.985	0.944
GENE-2	LECTURE-RANK	0.610	0.983	0.916
	CLASSICAL ML MODEL	0.247	0.965	0.788
	DIRECT	0.793	0.985	0.939
	LATENT	0.848	0.997	0.980
INNO-1	LECTURE-RANK	0.473	0.992	0.930
	CLASSICAL ML MODEL	0.569	0.992	0.936
	DIRECT	0.597	0.995	0.950
	LATENT	0.694	0.997	0.968
INNO-2	LECTURE-RANK	0.653	0.984	0.928
	CLASSICAL ML MODEL	0.644	0.984	0.928
	DIRECT	0.680	0.985	0.930
	LATENT	0.753	0.988	0.936
INNO-3	LECTURE-RANK	0.353	0.994	0.922
	CLASSICAL ML MODEL	0.141	0.986	0.792
	DIRECT	0.492	0.995	0.937
	LATENT	0.822	0.999	0.984

TABLE 9: Performance of LECTURE-RANK, DIRECT and LATENT models in predicting student survival

COURSE	MODEL	<i>start</i>	<i>mid</i>	<i>end</i>	<i>start-mid</i>
DISR	LECTURE-RANK	0.204	0.280	0.324	0.269
	DIRECT	0.304	0.400	0.470	0.372
	LATENT	0.417	0.454	0.629	0.451
WOMEN	LECTURE-RANK	0.538	0.518	0.415	0.533
	DIRECT	0.593	0.647	0.492	0.596
	LATENT	0.674	0.722	0.733	0.699
GENE-1	LECTURE-RANK	0.552	0.648	0.677	0.650
	DIRECT	0.647	0.755	0.784	0.692
	LATENT	0.705	0.755	0.789	0.778
GENE-2	LECTURE-RANK	0.449	0.431	0.232	0.699
	DIRECT	0.689	0.645	0.494	0.761
	LATENT	0.754	0.755	0.809	0.820
INNO-1	LECTURE-RANK	0.221	0.118	0.403	0.378
	DIRECT	0.383	0.304	0.846	0.692
	LATENT	0.571	0.460	0.854	0.778
INNO-2	LECTURE-RANK	0.232	0.464	0.456	0.301
	DIRECT	0.438	0.600	0.637	0.565
	LATENT	0.605	0.676	0.794	0.648
INNO-3	LECTURE-RANK	0.104	0.188	0.203	0.113
	DIRECT	0.202	0.405	0.478	0.293
	LATENT	0.309	0.574	0.803	0.428

TABLE 10: Early prediction performance of LECTURE-RANK, DIRECT and LATENT models in time-periods *start*, *mid*, *end*, and *start-mid*

into three equal parts according to the duration of the course: *start*, *mid*, and *end*. We combine *start* and *mid* time periods to get data till *mid* part of the course, which we refer to as *start-mid*. *start-end* refers to data collected over the entire course. In all, we consider five time-periods in our experiments: *start*, *mid*, *end*, *start-mid*, and *start-end*. The student survival labels are the same as for the complete dataset (i.e., whether the student submitted the final quizzes/assignments at the end of the course), but our models are only given access to data from the early parts of the

course. All features are re-calculated to include data from only the specific time period in consideration. For example, POSTS(U,P) is modified to only include posts in that specific time period.

Table 10 lists the performance metrics for our two models using different splits in the data. Similar to the results in Table 9, the change in the AUC-PR (Neg.) scores are negligible and close to optimal for all models because of class imbalance. To highlight the strength our models, we only report the AUC-PR (Pos.) scores of the models. Early prediction scores under *start*, *mid*, and *start-mid* indicate that our model can indeed make early survival predictions reliably. As the data available is closer to the end of the course, models make better predictions. Similar to the previous experimental setting, the LATENT model achieves the highest prediction quality. We observe that the LATENT model consistently outperforms the DIRECT model on all time periods across seven courses. The LATENT model also significantly outperforms the DIRECT model in the *start* time period, making it a very useful tool for instructors to predict student survival early on in the course.

From the results, it appears that the middle phase (*mid*) is the most important phase to monitor student activity for predicting whether the student will survive the length of the course. Our model produces higher AUC-PR values when using data from the *mid* phase, compared to the settings where we use data from the *start* phase, and an almost equal value when compared to *start-mid*. We hypothesize that this is due to the presence of a larger student population in the *start* phase that fails to remain engaged until the end. This phenomenon is typical in both traditional and online classrooms where students familiarize themselves with the course and then decide whether to stay or drop out. Eliminating data collected from this population helps improve our prediction of student survival, as indicated by an increase in performance values for *mid*.

5.5 Survival Prediction on Unseen Courses

So far, we demonstrated the predictive ability of our models in predicting survival on courses by training on data from the same course. But for new courses which haven't yet accumulated performance and survival data for students, it is not possible to train on data from the same iteration of the course. Models trained on other courses, but having good predictive power in predicting student success on new or previously *unseen* courses will be very beneficial. Predicting student survival on courses in progress helps instructors monitor and track student engagement and initiate interventions promptly before students disengage and dropout. We demonstrate the extensibility of our models in predicting survival on new courses by training on data from one course and testing on a different course.

Table 11 gives the performance metrics for DIRECT and LATENT models, training on the course indicated by TRAIN COURSE and testing on data from TEST COURSE. The scores indicate that both our models can predict survival on new courses reliably. We experiment on two different combinations of train and test courses: i) the train and test courses are drawn from different iterations of the same course, ii) the train and test courses are drawn from different courses. For example, the first three rows in Table 11 provide results for training on a different iteration of the same course. The last two rows give results for training on INNO-1 and testing on GENE-2, and training on GENE-2 and testing on INNO-2, respectively. The second experiment is especially helpful for predicting survival in new courses, which do not have any

previous iterations to train on. In both these cases, we observe that our models achieve good predictive performance comparable to training on the same course.

TRAIN	TEST	MODEL	AUC-PR Pos.	AUC-PR Neg.	AUC-ROC
INNO-1	INNO-2	DIRECT	0.721	0.989	0.945
		LATENT	0.713	0.987	0.933
INNO-1	INNO-3	DIRECT	0.506	0.996	0.940
		LATENT	0.719	0.998	0.978
GENE-1	GENE-2	DIRECT	0.737	0.987	0.934
		LATENT	0.762	0.995	0.962
INNO-1	GENE-2	DIRECT	0.709	0.986	0.932
		LATENT	0.853	0.997	0.979
GENE-2	INNO-2	DIRECT	0.723	0.990	0.945
		LATENT	0.683	0.985	0.922

TABLE 11: Prediction performance of DIRECT and LATENT models in training on one course and testing on another course.

TRAIN COURSE	TRAIN TIME PERIOD	TEST COURSE	TEST TIME PERIOD	MODEL	AUC-PR
					Pos.
INNO-1	<i>start-end</i>	INNO-2	<i>start</i>	DIRECT	0.628
				LATENT	0.658
INNO-1	<i>start</i>	INNO-2	<i>start</i>	DIRECT	0.618
				LATENT	0.652
INNO-1	<i>start-end</i>	INNO-3	<i>start</i>	DIRECT	0.318
				LATENT	0.400
INNO-1	<i>start</i>	INNO-3	<i>start</i>	DIRECT	0.363
				LATENT	0.394
INNO-1	<i>start-end</i>	GENE-2	<i>start</i>	DIRECT	0.712
				LATENT	0.885
GENE-2	<i>start-end</i>	INNO-2	<i>start</i>	DIRECT	0.625
				LATENT	0.657
GENE-2	<i>start</i>	INNO-2	<i>start</i>	DIRECT	0.627
				LATENT	0.657


TABLE 12: Early prediction performance of DIRECT and LATENT models in training on one course and testing on another course.

5.6 Early Survival Prediction on Unseen Courses


Next, we investigate the reliability of our models in early prediction when they are trained on data from a different course. Achieving good early prediction performance is especially helpful to courses in progress, allowing instructors to intervene before the students disengage and dropout. Here, we consider four different experiment settings, to understand the capabilities of our models when trained on different training data sets. We first consider the two experiment settings that we considered in Section 5.5: i) the train and test courses are drawn from different iterations of the same course, ii) the train and test courses are drawn from different courses. For each of these settings, we consider two possible variations on the training dataset: i) training on data from an entire course different from the test course (indicated by *start-end*), and ii) training on data from the time-period corresponding to the time-period of the test course. Hence, in all, we consider four different combinations of train and test datasets. We evaluate the prediction performance on the most challenging early prediction period *start*, as this time period has the least amount of data. Table 12 gives the early prediction results. Notice that both our models achieve good prediction performance, with the LATENT model performing better than the DIRECT model in most cases. We observe that training on data from different iteration of the same course often yields

better prediction performance than training on the data from the same iteration of the course (comparing results for time period *start* in Table 10 and Table 12), which demonstrates the utility of our models across iterations of the same course. We observe that training on entire data from a different course is better than training on the exact time period (indicated by *start*), indicating our models can potentially be trained on existing courses and used in earlier time periods of new courses to facilitate interventions.


5.7 Feature Analysis




(a) Feature analysis in DISR course



(b) Feature analysis in WOMEN course




(c) Feature analysis in GENE-1 course




(d) Feature analysis in GENE-2 course

Fig. 3: Bar graph showing AUC-PR (Pos.) value upon removal of each feature from the DIRECT model across time periods


Here, we perform a comprehensive feature analysis to understand the predictability of each feature in predicting student success in online courses. We group the features into sets of features: a) *post*: features related to posting in forums, including linguistic and structural features derived from forum posts, b) *view*: viewing forum content, c) *lecture*: viewing lectures and taking associated quizzes, d) *temporal*: temporal features, and e) *all*: the entire model with all the features. We evaluate the contribution of each feature group in predicting student success, by



(a) Feature analysis in INNO-1 course



(b) Feature analysis in INNO-2 course



(c) Feature analysis in INNO-3 course

Fig. 4: Bar graph showing AUC-PR (Pos.) value upon removal of each feature from the DIRECT model across time periods

leaving each feature group out and observing the resulting change in the area under precision-recall curve and area under ROC values. To do so, we omit all PSL rules that mention the feature group. For example, to evaluate the importance of the first feature group *post*, we remove all features related to posting in forums such as POST-ACTIVITY, POSTS, POLARITY, and structural rules connecting forum posts. Feature groups have varying levels of predictability across the different time periods. We compare the predictability of the feature groups across the five time periods discuss in Section 5.4: *start*, *mid*, *end*, *start-mid*, and *start-end*. Figure 3 and 4 plots the results from the experiments removing each feature-group across the different time periods. The decrease in value from *all* corresponds to the importance of each feature group in the model.

From Figures 3 and 4, we observe that the *lecture* feature group is consistently important for predicting student survival, indicating that it is the most prevalent form of interaction of MOOC participants on the MOOC website. This is especially evident in the *mid* and *end* phases, where *lecture* is a very important feature. In some courses, it is a very strong feature from the *start* phase (DISR, WOMEN, GENE-1, and GENE-2) (Figure 3), while in the INNO courses (4), it only becomes relevant in the *mid* and *end* phases. Discussion forums serve as a platform connecting students worldwide enrolled in the course, hence activity in the discussion forums also turns out to be a strongly contributing feature. Since, the concentration of forum posts in the courses analyzed is more in the *mid* and *end* phases, posting in forums is accordingly more important during the *mid* and *end* phases. Also, in the *start* phase of the course, most posts are about students

introducing themselves and getting to know other people enrolled in the course. These posts are not very predictive of student engagement and their subsequent performance or survival in the course. Simply viewing content on the forums (*view*) is also a strong feature, contributing consistently in all phases across all courses. In fact, from Figures 3 and 4, we can see that the feature strength of forum views is second only to lecture views. We also observe that the effect of lecture viewing is less significant in some courses, while forum viewing is more significant instead (WOMEN, GENE-2, and INNO-3). This can be attributed to the presence of active discussions encouraged in the course by the instructor, starting discussion topics where many students participating. A larger fraction of students view these posts and use them to understand the material, hence forum viewing in these courses has a significant impact on performance. This further ascertains the importance of passive engagement in online courses. *Temporal* features are a strong feature in the early part of the course, particularly in the *start* phase across all seven courses. But, they decline as a predictive feature in the *mid* and *end* phases. The data suggests that this is due to the larger volume of students dropping out in the early part of the course, making it an excellent predictor for student survival in the *start* phase. As the student population grows steady, *temporal* features start to decline as a predictive feature.


We observe a similar trend when we observe the weights of the rules in our DIRECT and LATENT models. We observe that the rules containing features from the lecture feature-group obtain the highest learned weights. This is followed by rules containing the *view* feature group. Following this, in the latent model, are rules containing ENGAGEMENT-PASSIVE, which is followed by rules containing ENGAGEMENT-ACTIVE. From this we note that ENGAGEMENT-PASSIVE is more predictive of student success than ENGAGEMENT-ACTIVE, which conforms to the observations in the classroom settings. The next prominent set of rules are rules containing the *post* feature group. This is followed by rules containing the temporal features in early time periods. Rules containing all other features come after the rules mentioned above.

5.8 Gaining Insight from Latent Engagement Assignments


So far, we demonstrated the utility of the latent engagement variables in performance prediction. Going beyond measuring the impact of engagement on performance prediction, we are interested in understanding the value of the engagement information our model uncovers.

In this section, we further dissect the latent engagement values to see how student engagement evolves as the course progresses. We track the changes in engagement assignments patterns for several interesting student populations and discuss potential explanations for these changes. We categorize students that drop out of the course according to the time period in which they dropped out. We analyze the student engagement values predicted by the model for three groups of students—(1) students dropping out in the *mid* phase, (2) students dropping out in the *end* phase, and (3) students continuing until course completion.


We train our models on data from *start*, *mid*, and *end* phases of the course and record the engagement values for the students in these three periods. We consider three groups of students: 1) students dropping out in the *mid* phase, 2) students dropping out in the *end* phase, and 3) students continuing till the end.



(a) Engagement patterns of students that dropped out of the class in the middle phase



(b) Engagement patterns of students that dropped out of the class in the end phase



(c) Engagement patterns of students that survived the complete class

Fig. 5: Bar-graph showing the distribution of engagement label assignments at three time points throughout the class. We capture engagement transition patterns by coloring the bars according to the engagement assignments of students at the previous time point.

Students dropping out in the *mid* phase stop participating in course activities sometime during middle phase. Similarly, students dropping out in the *end* phase stop participating in the course sometime during the *end* phase. The students are classified into one of the engagement types by considering the dominant value of engagement as predicted by the model. Using this we distinguish between the different engagement types for different populations of students and uncover their movement from one engagement type to another and understand how engagement-mobility patterns relate to student survival.

Figure 5 describes the student engagement values predicted by the model for the three classes of students. For each student group, we provide a bar graph, showing the different engagement assignment levels at each time span (*start*, *middle*, *end*). The labels D, EA and EP refer to values for latent variables DISENGAGEMENT, ACTIVE-ENGAGEMENT and PASSIVE-ENGAGEMENT, respectively. Let us first consider Figure 5(a). In the start period, we first categorize students into three forms of engagement D, EA, and EP, respectively. The three engagement types are denoted by the colors *red*, *yellow*, and *green*, respectively in the start period. In the middle period, we capture the total number students in each engagement category in the columns D, EA, and EP. In order to track student engagement patterns, we color code the bars in

the middle and end phases according to the previous engagement assignments of the students, with the colors *red*, *yellow*, and *green* capturing the number of students with engagement type DISENGAGEMENT, ACTIVE-ENGAGEMENT, and PASSIVE-ENGAGEMENT in the previous time period, respectively. Each bar therefore consists of the combination of three smaller bars, colored differently, capturing the previous engagement values.

In Figure 5(a), in the middle phase, there is almost equal percentage of students moving from DISENGAGEMENT, ACTIVE-ENGAGEMENT, and PASSIVE-ENGAGEMENT in the start phase. EA students start to move toward disengagement in the middle phase. While some EP students, who are not taking quizzes in middle phase, still follow the course passively, placing them in EP rather than D. *We hypothesize that these students may be more likely to respond to intervention than the already disengaged students.* In Figure 5(b), it can be seen that, out of the students that drop out eventually in the *end* phase, about half of them are in EP. Finally, Figure 5(c) suggests that most engaged students only exhibit passive forms of engagement in the *start* and *mid* phases of the course. While in the *end* phase, students tend to become more actively engaged in the course. All these results corroborate the importance of taking into account passive engagement. Several education works state the importance of passive forms of engagement and their subtlety [8], [9], [10], [11]. With our thorough construction of features contributing to passive engagement, we are able to observe similar trends in the online setting. In all these classes of students, passive engagement is a more prevalent type of engagement than active, stressing the fact that careful observation of passive engagement (which includes subtle activities such as viewing forum posts) can help MOOC instructors assess student health.

6 CONCLUSION

In this work, we take a step toward helping MOOC instructors and optimizing experience for MOOC participants by modeling latent student engagement using data-driven methods. We formalize, using HL-MRFs, that student engagement can be modeled as a complex interaction of behavioral, linguistic and social cues, and we model student engagement types as latent variables over these cues. We demonstrate the effectiveness and reliability of our models through a series of experiments across seven MOOCs from different disciplines, analyzing their predictive performance on predicting student success, early prediction of student survival, survival prediction on unseen courses, and a detailed feature analysis capturing the contribution of each feature group in predicting student success. Our models construct interpretations for latent engagement variables from data and predict student course success indicators reliably, even at early stages in the course, particularly on previously unseen courses, making them very useful for instructors to assess student engagement levels. These results are a first step toward facilitating instructors' intervention at critical points for courses in progress, thus helping improve course retention rates. The latent formulation we present can be extended to more sophisticated modeling by including additional latent factors that affect academic performance such as motivation, self-regulation and tenacity. Our models can also be integrated into an automatic framework for monitoring student progress and initiating instructor interventions. These compelling directions for future interdisciplinary investigation can provide a better understanding of MOOC students.

REFERENCES

[1] P. Brusilovsky and E. Millán, "User models for adaptive hypermedia and adaptive educational systems," in *Lecture Notes In Computer Science*, 2007, vol. 4321, pp. 3–53.

[2] A. Ramesh, D. Goldwasser, B. Huang, H. Daume III, and L. Getoor, "Learning latent engagement patterns of students in online courses," in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2014.

[3] —, "Modeling learner engagement in MOOCs using probabilistic soft logic," in *NIPS Workshop on Data Driven Education*, 2013.

[4] M. G. Postel, H. A. de Haan, E. D. Ter Huurne, E. S. Becker, and C. A. de Jong, "Effectiveness of a web-based intervention for problem drinkers and reasons for dropout: randomized controlled trial," *Journal of Medical Internet Research (JMIR)*, vol. 12, no. 4, p. e68, 2010.

[5] V. Tinto, *Leaving college: Rethinking the causes and cures of student attrition*. ERIC, 1987.

[6] L. S. Hawken and R. H. Horner, "Evaluation of a targeted intervention within a schoolwide system of behavior support," *Journal of Behavioral Education*, vol. 12, no. 3, pp. 225–240, 2003.

[7] R. McWilliam, C. M. Trivette, and C. J. Dunst, "Behavior engagement as a measure of the efficacy of early intervention," *Analysis and Intervention in Developmental Disabilities*, vol. 5, no. 1-2, pp. 59–71, 1985.

[8] K. J. Herrmann, "The impact of cooperative learning on student engagement: Results from an intervention," *Active Learning in Higher Education*, vol. 14, no. 3, pp. 175–187, 2013.

[9] K. A. Rocca, "Student participation in the college classroom: An extended multidisciplinary literature review," *Communication Education*, vol. 59, no. 2, pp. 185–213, 2010.

[10] G. D. Kuh, "What we are learning about student engagement from nsse: Benchmarks for effective educational practices," *Change: The Magazine of Higher Learning*, 2003.

[11] R. M. Carini, G. D. Kuh, and S. P. Klein, "Student engagement and student learning: Testing the linkages," *Research in Higher Education*, vol. 47, no. 1, 2006.

[12] S. J. Richards, "A handbook of parametric survival models for actuarial use," *Scandinavian Actuarial Journal*, vol. 2012, no. 4, pp. 233–257, 2012.

[13] R. Crues, N. Bosch, M. Perry, L. Angrave, N. Shaik, and S. Bhat, "Refocusing the lens on engagement in MOOCs," in *Proceedings of the ACM Conference on Learning at Scale (L@S)*, 2018.

[14] R. F. Kizilcec, C. Piech, and E. Schneider, "Deconstructing disengagement: Analyzing learner subpopulations in massive open online courses," in *Proceedings of the International Conference on Learning Analytics and Knowledge (LAK)*, 2013.

[15] A. Anderson, D. Huttenlocher, J. Kleinberg, and J. Leskovec, "Engaging with massive online courses," in *Proceedings of the International Conference on World Wide Web (WWW)*, 2014.

[16] J. Qiu, J. Tang, T. X. Liu, J. Gong, C. Zhang, Q. Zhang, and Y. Xue, "Modeling and predicting learning behavior in MOOCs," in *Proceedings of the International Conference on Web Search and Data Mining (WSDM)*, 2016.

[17] A. S. Lan, C. G. Brinton, T.-Y. Yang, and M. Chiang, "Behavior-based latent variable model for learner engagement," in *Proceedings of the International Conference on Educational Data Mining (EDM)*, 2014.

[18] J. Lui and H. Li, "Exploring the relationship between student pre-knowledge and engagement in MOOCs using polytomous IRT," in *Proceedings of the International Conference on Educational Data Mining (EDM)*, 2017.

[19] J. He, J. Bailey, B. I. Rubinstein, and R. Zhang, "Identifying at-risk students in massive open online courses," in *Proceedings of the Conference on Artificial Intelligence (AAAI)*, 2015.

[20] J. Gardner and C. Brooks, "A statistical framework for predictive model evaluation in MOOCs," in *Proceedings of the ACM Conference on Learning@ Scale (L@S)*, 2017, pp. 269–272.

[21] S. Kotsiantis, C. Pierrakeas, and P. Pintelas, "Preventing student dropout in distance learning using machine learning techniques," in *Knowledge-Based Intelligent Information and Engineering Systems*, ser. Lecture Notes in Computer Science, V. Palade, R. Howlett, and L. Jain, Eds., 2003, vol. 2774, pp. 267–274.

[22] D. Clow, "MOOCs and the funnel of participation," in *Proceedings of the International Conference on Learning Analytics and Knowledge (LAK)*, 2013.

[23] G. Balakrishnan, "Predicting student retention in massive open online courses using hidden Markov models," Master's thesis, EECS Department, University of California, Berkeley, 2013.

[24] D. Yang, T. Sinha, D. Adamson, and R. Penstein, "Turn on, tune in, drop out: Anticipating student dropouts in massive open online courses," in *NIPS Workshop on Data Driven Education*, 2013.

[25] A. Cernezel, S. Karakatic, B. Brumen, and V. Podgorelec, "Predicting grades based on students' online course activities," in *Knowledge Management in Organizations (KMO)*, 2014, pp. 108–117.

[26] P. J. Guo, J. Kim, and R. Rubin, "How video production affects student engagement: An empirical study of MOOC videos," in *Proceedings of the ACM Conference on Learning @ Scale Conference (L@S)*, 2014.

[27] I. Nawrot and A. Doucet, "Building engagement for MOOC students: Introducing support for time management on online learning platforms," in *Proceedings of the International Conference on World Wide Web (WWW)*, 2014.

[28] M. Wen, D. Yang, and C. P. Rosé, "Linguistic reflections of student engagement in massive open online courses," in *Proceedings of the International Conference on Web logs and Social Media (ICWSM)*, 2014.

[29] A. Ramesh, S. Kumar, J. Foulds, and L. Getoor, "Weakly supervised models of aspect-sentiment for online course discussion forums," in *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2015.

[30] S. Tomkins, A. Ramesh, and L. Getoor, "Predicting post-test performance from online student behavior: A high school mooc case study," in *International Conference on Educational Data Mining (EDM)*, 2016.

[31] W. Chen, C. G. Brinton, D. Cao, A. Mason-singh, C. Lu, and M. Chiang, "Early detection prediction of learning outcomes in online short-courses via learning behaviors," *IEEE Transactions on Learning Technologies*, 2018.

[32] S. H. Bach, M. Broecheler, B. Huang, and L. Getoor, "Hinge-loss markov random fields and probabilistic soft logic," *Journal Machine Learning Research*, vol. 18, no. 1, pp. 3846–3912, Jan. 2017.

[33] S. H. Bach, B. Huang, and L. Getoor, "Learning latent groups with hinge-loss Markov random fields," in *ICML Workshop on Inferring: Interactions between Inference and Learning*, 2013.

[34] T. Wilson, P. Hoffmann, S. Somasundaran, J. Kessler, J. Wiebe, Y. Choi, C. Cardie, E. Riloff, and S. Patwardhan, "OpinionFinder: A system for subjectivity analysis," in *Proceedings of HLT/EMNLP on Interactive Demonstrations*, 2005.

[35] A. S. Lan, J. C. Spencer, Z. Chen, C. G. Brinton, and M. Chiang, "Personalized thread recommendation for MOOC discussion forums," *CoRR*, vol. abs/1806.08468, 2018.



Arti Ramesh is currently an assistant professor at State University of New York, Binghamton. She received her PhD in Computer Science from University of Maryland, College Park.



Dan Goldwasser is an assistant professor at Purdue University. Before that, he was a post-doctoral associate at University of Maryland, College Park. He received his Ph.D. degree in computer science from University of Illinois at Urbana-Champaign.



Bert Huang is an assistant professor at Virginia Tech. Before that, he was a postdoctoral associate at University of Maryland, College Park. He received his Ph.D. degree from Columbia University.



Hal Daume III is an associate professor in Computer Science at the University of Maryland, College Park. He was previously an assistant professor in the School of Computing at the University of Utah. He earned his PhD at the University of Southern California with a thesis on structured prediction for language.



Lise Getoor is a professor in the Computer Science Department at UC Santa Cruz and the director of the UC Santa Cruz D3 Data Science Center. She received her PhD from Stanford University, her MS from UC Berkeley, and her BS from UC Santa Barbara, and was a professor at the University of Maryland, College Park from 2001-2013.