

Using Discussion Board Data to Hire Teaching Assistants

Arnaud Deza
University of Toronto
St. George
Toronto, Canada
arnaud.deza@mail.utoronto.ca

Haocheng Hu
University of Toronto
St. George
Toronto, Canada
hc.hu@mail.utoronto.ca

Vaishvik Maisuria
University of Toronto
Mississauga
Mississauga, Canada
vaishvik.maisuria@mail.utoronto.ca

Michael Liut^{*}
University of Toronto
Mississauga
Mississauga, Canada
michael.liut@utoronto.ca

Andrew Petersen^{*}
University of Toronto
Mississauga
Mississauga, Canada
andrew.petersen@utoronto.ca

Bogdan Simion^{*}
University of Toronto
Mississauga
Mississauga, Canada
bogdan.simion@utoronto.ca

ABSTRACT

Teaching Assistants (TAs) fill critical roles in instructional teams, spending one-on-one time with students, providing feedback, and serving a critical mentoring role. As a result, the process of selecting TAs is important, yet instructors often have little information when making hiring decisions. We used data collected from a course discussion board to identify statistics that are indicative of participation and that might be useful for hiring TAs. We then analyzed opast TA hires and found that there are opportunities to use discussion board metrics to increase the number of potentially qualified applicants for TA roles and to better inform instructors about the engagement of past students as they select candidates.

Author Keywords

Educational data mining; learning analytics; discussion board; discussion board analytics

CCS Concepts

•**Social and professional topics** → **Computing education; Computer science education; Applied computing** → **Learning management systems; Interactive learning environments; Collaborative learning;**

INTRODUCTION

Large courses require Teaching Assistants (TAs) to operate, and in larger institutions, many of those TAs are former students or current undergraduate students. TAs are important members of the instructional team, with many spending significant time mentoring and providing feedback. As a result, the process of selecting TAs is an important, if under-studied, element of course preparation.

The TA hiring process varies by institution, but in many cases, the hiring instructor has relatively little information about applicants who have not held a TA position before. Most applicants will not have relevant experience, making performance

in related courses the most commonly available piece of evidence. Instructors faced with such hiring decisions often have no information about students beyond these grades. In particular, instructors may not have any information about an applicant's level of engagement with their peers or about their communication skills, yet we know that communication skills are necessary for effective teaching behaviours [3]. Data from discussion boards can play an effective role in bridging this knowledge gap in contexts in which former students (in particular, current undergraduate students) might be considered for TA positions. In our particular context, undergraduate TAs play a crucial part by filling peer-mentoring roles in first year courses while they are senior undergraduates.

In this study, we use data collected on the Piazza discussion board [10] and analyze which participation statistics (e.g., posts viewed, number of days active, number of questions asked and answered, etc.) can be utilized to identify additional potential candidates for TA positions. These indicators are mapped against each respective student's course grade to provide a better overview of a student. Additionally, by cross-referencing with hiring data from a subsequent year:

1. We find that while several applicants with high grades and high active participation/basic engagement in the prior year do get hired, there are some who are not even considered. While interviewing hiring instructors it was found that this was due to a lack of information about those candidates. We expect that these analytics may benefit instructors to make more informed decisions in selecting and interviewing TAs who have the right attributes.
2. We find that quite a few students with high grades, who were active participants on the discussion board in the prior year, simply do not apply to become a TA at all. This suggests that better communication about what it takes to be a TA and how to apply might attract more applicants and increase the pool of good candidates.

In the next section, we look at some of the learning theory that suggests that facilitating community is an important as-

^{*}corresponding authors

pect of teaching, to motivate why the data we're suggesting to use might be relevant. We also briefly discuss other work that has used discussion board data to identify leaders in online communities. In the Methods section, we describe how the data was collected, how we identified heuristics to use, and what threats to validity we have identified. The Data and Discussion sections present the data collected and discuss how that data might change our hiring practices. Finally, we summarize and identify future directions for work.

RELATED WORK

Multiple papers have found that engagement in discussions – one aspect of social engagement in a course – is correlated with higher performance [14, 2], and more broadly, that interaction between students is a key driver of performance [11]. These results reflect Lave and Wenger's theory of situated learning [8], which posits that acquisition of professional skills is facilitated by social engagement and participation in a community, and are also related to Kearsley and Schneiderman's engagement theory, which suggests that learning requires meaningful engagement in activities, including interactions with peers [6].

These results explain why the creation of an online learning community is such an important part of many classrooms – and particularly, in online and hybrid course environments. Epp et al. found that instructor facilitation has a particularly important place in the creation of community [4]. Therefore, in large courses, where the instructor works with a team of teaching assistants to deliver the course, the ability of the teaching assistants to mentor and facilitate community in an online environment is important and should be considered when building a teaching team. Preece and Shneiderman have found that cultivating leaders in online environments requires that both mentoring activities and existing leaders be recognized [12], so hiring (effectively, promoting) active participants in online communities has the additional benefit of promoting leadership activity in the community.

Several groups have mined discussion board activity to identify leaders in online communities [2, 5]. Huffaker used social network and text analysis to examine the behaviour of leaders in online communities [5]. More recently and in the field of computing education, Desai et al. also used social network analysis to identify key figures in online student discussions [2]. Furthermore, in professional online settings, leaders have more social capital and are identified by their peers based on the amount of relevant knowledge they can bring to bear from their professional lives [13].

Despite these successes, collecting the right data to accurately and comprehensively identify which participants are particularly engaged and which are demonstrating leadership activities may be challenging. Social connections on discussion boards tend to be sparse [14]. Furthermore, we may not, typically, be evaluating interactions with sufficient depth. Oztok et al. took a social capital theory perspective while examining discussion board interactions and found that current models of social interactions do not sufficiently consider the quality of interactions and the density of links in the social network [9]. Nevertheless, even fairly lightweight data can identify

the most active contributors, which may be enough to better support the TA hiring process.

The discussion board we used to collect data was Piazza [10]. Several other groups have used Piazza discussion board data to predict students at risk [7] or to extend the functionality of the board to make it more responsive to student queries [1]. However, we are not aware of other work that uses data from Piazza – or another discussion board – to identify engagement behaviours for the purpose of informing hiring.

METHODS

In this section, we first identify the data available and then discuss constraints and preferences on TA hiring that speak to specific attributes of the data available. Finally, we identify a set of threats to validity that should be considered prior to reviewing the data and analysis.

Data and Data Cleaning

The data used in this investigation are the last two years of discussion board and course performance data in a CS1 course (2 separate offerings of approximately 1,000 students per offering, which we refer to further as CS1 Year 1 and CS1 Year 2) and a CS2 course (2 separate offerings of approximately 750 students per offering, which we refer to as CS2 Year 1 and CS2 Year 2), where the CS1 course is as a prerequisite. All students in each of the courses were automatically enrolled in their respective discussion board, but they were not required to participate. Discussion board data was collected, students were linked with their respective grades in the course and then de-identified for analysis.

The de-identification process ensured that students were provided with a unique identifier that cannot be linked back to them if the dataset were leaked. This identifier did allow us to link their grades, TA status (did they apply/were they hired), and the discussion board data.

The discussion board used in both courses was Piazza [10]. Student aggregate data collected through Piazza includes the number of days online, posts made, and posts viewed. Additional data, such as the number of editors to any one post, the number of students who have marked a post good/helpful or have had their posts marked as such, the category of the post, etc., require further mining and processing.

Individual data collected through Piazza includes the classification of the post (question or answer, original post or follow-up, etc.), the general category which the post is associated to, who created the post, the subject line and content of the post, the date/time it was created, if the students opted to remain anonymous to their peers, and whether a post was endorsed by an Instructor/Teaching Assistant. Student grade data includes a detailed breakdown of each piece of evaluation. This includes their midterms/tests, final exam, labs, exercises, and assignments, as well as their final course grade.

As students in CS1 Year 1 and CS2 Year 1 would be eligible for TA hiring, at the earliest, in either CS1 Year 2 or CS2 Year 2, we collected information on all second year TA applicants for the latter two offerings to identify which second year applicants applied and which were successful in getting

a position. Additionally, CS1 Year 2 and CS2 Year 2 TAs are eligible for TA hiring in the Year 3 offering of CS1 and CS2, so we collected information from that hiring round as well.

Hiring Practices and Hypotheses

When hiring Teaching Assistants (TAs), instructors typically look for content expertise and ability to interact with students in guiding them through the learning process (whether through help with practical tasks, steering questions on a discussion board, etc.). To some extent, instructors may be looking for programming maturity, so it is perhaps not unexpected that an upper-year student or graduate student would be preferred to a second-year student applicant. In addition, certain academic institutions may impose requirements on TA hiring (e.g. minimum grade, union restrictions, graduate student preferences, etc.) which directly impact an instructor's decision. However, in many cases the pool of qualified senior undergraduates and/or available graduate students is insufficient for the number of undergraduate courses that must be staffed, so younger students may need to be hired, or existing candidates may be asked to take on additional TA sections.

Instructors do not typically have information about a candidate's suitability for a TA position beyond grades, which speaks only to the technical ability of the candidate. Having evidence of active engagement in discussions and helping their peers in prior courses could add valuable insight when selecting candidates for interviews or when making a hiring decision between two academically strong candidates.

We inspected all Piazza participation statistics and classify them into two aggregate indicators of engagement/participation on the discussion board:

1. The number of days active and posts viewed indicates a basic level of engagement where the student is at least following the activity on the discussion board. As these two indicators are somewhat intertwined, we take the aggregate of posts viewed and days active as the first indicator of engagement. We name this the *basic engagement* indicator.
2. The number of questions and answers indicate a more active engagement beyond passively keeping up with the course. While this indicator is not perfect (i.e., not all questions are high quality, nor is every answer), it shows active involvement to ask for clarification or to help peers by answering questions or follow-ups. We use this as a second indicator of engagement and label it as the *active participation* indicator.

We hypothesize that when the above engagement indicators are matched with course grades, we can identify potentially suitable TAs and provide such information to course instructors for further evaluation as they decide who to interview and hire. We expand on this in the latter half of the Data section.

Threats to Validity

We have identified two major threats in our approach: small sample sizes and evidence of engagement that cannot be captured from the discussion board statistics.

First, as mentioned in the previous section, hiring practices at our institution prioritize graduate students or upper-year undergraduates rather than second-year students. As a result, the number of TA positions available to second-year students tends to be small. For instance, for the CS1 Year 1 offering, only 3 out of 25 TAs (12%) were second-year undergraduate students, while for CS1 Year 2, we had 2 out of 21 (10%) TA hires for CS1 and 5 out of 36 (14%) hires for CS2.

Second, although discussion board indicators contain valuable insight, they are not the only evidence of engagement in helping peers. For instance, some students may be actively participating in class discussions or lab activities, but may not register as high activity in an online forum. As a result, our effort to connect applicant data to the kind of engagement we want from a prospective TA is limited. We can only include factors that can be identified exclusively from discussion board behaviour, and these are not solely indicative of a successful student nor a potentially successful TA. As we will see later, although some high-performing students are very active on Piazza, other high-performing students are not, and many lower-performing students are equally active.

Third, the students were not aware, when this data was collected, that their engagement metrics could be used to identify them as a potential TA candidate, and if this data were to be actually used for hiring, we would need to inform them. Their behaviour on the discussion board – and as a result, the metrics we calculate – might have been different had they known that their data would be used. This issue is explored further in the Discussion Section.

DATA

Table 2 contains basic statistics for each course: number of students who were still enrolled by the end of the corresponding term, number of active Piazza users (including instructional staff, removing those who have viewed no posts), average number of days active per user, average number of posts viewed per user, average number of posts made per user, average number of contributions made per user, and the total number of contributions made on the discussion board (this includes posts and follow-up discussions). We identified indicators of “basic engagement” as discussed in the Hiring Practices section: being active on a day or reading a post or follow-up; those are strongly correlated. Similarly, we identified indicators of “active participation”: posting and answering are also correlated. The others are not. Bonferroni correction was performed due to the large number of statistical tests performed. See Table 1 for correlations.

	Basic Engagement	Active Participation	Grades and Basic Engagement	Grades and Active Participation
Pearson Correlation	0.70	0.34	0.37	0.18
Spearman Correlation	0.81	0.50	0.44	0.32
Kendall Correlation	0.63	0.44	0.31	0.24

Table 1. Correlation of Data ($p < 0.001$ for all entries)

	CS1 Year 1	CS1 Year2	CS2 Year1	CS2 Year2
Final Student Count	948	942	597	651
Number of Active Piazza Users	1,011	997	648	694
Average Days Active	84.8	66.8	83.3	61.8
Average Posts Viewed	236.43	225.8	325.5	571.0
Average Posts Made	1.06	1.9	2.3	4.4
Average Contributions	5.1	9.4	11.0	27.1
Total Contributions Made	5,218	11,341	7,607	19,215

Table 2. Course Data

The correlations between engagement and grades in Table 1 are low. To explore this further, we plotted student course grades against the *basic engagement* and *active participation* indicators. The results are in Figures 1 and 2. We see that while high engagement is generally (though not universally) indicative of high performance, low engagement is not indicative of low performance. There are many reasons not to be engaged, but most people who are highly engaged are doing well. In those figures, we also cross-reference against hiring information from CS1 Year 2 and CS2 Year 2 and further annotate the figures to differentiate between students who successfully obtained a TA position (blue squares), those who applied to be a TA but were not hired (red crosses), and students who did not apply to be a TA (grey circles). Those hired are always academically strong, but they are not universally active according to our engagement measures.

Finally, we attempted to identify potential current TA candidates that would have been overlooked if the instructors were not provided with discussion board analytics as a supplementary factor. After making their normal hiring decisions, the instructors were provided with two sets of discussion board data (those plotted in the X-axis of Figures 1 and 2) and were asked to reconsider their TA selections with the additional discussion board analytics. Instructors provided a list of students they would “reconsider” offering a position given this extra data; those students are plotted with light blue triangles in the figures with Year 2 (current year) data.

DISCUSSION

We do not believe there is an effective, generally useful formula for identifying good TAs. Given that each academic institution and course may have its own requirements for hiring (see the Hiring Practices section) there is no real way to generate a universal formula that describes what a good TA looks like, let alone one that is “hirable” under institution rules. Instead, figures 1 and 2 suggest that there is the potential to use discussion board data to identify additional potentially-qualified applicants who would otherwise not be considered.

When given additional discussion board information, the instructors consulted expressed interest in student applicants

with a combination high engagement and grades. They agreed that there may have been good TA applicants that were overlooked due to a “lack of information about the applicant(s)” and, when provided with discussion board data, identified a set of candidates they would have considered. If instructors had this information on hand, they would have interviewed more applicants who have a record of consistent engagement. However, instructors also stated that the people they had actually hired had been effective TAs, which provides further reinforcement that this information should not be used to assist instructors in providing additional choices rather than being taken as an indicator of “better” choices.

Beyond improving the data available when making hiring decisions, we believe this data can be useful to identify potentially qualified applicants who did not apply. In every figure, there are students who did not apply to be TAs (grey circles) who had higher engagement, higher grades, or both when compared to TAs who were actually hired. This indicates potential missed opportunities for qualified TAs. While we cannot determine why these students did not choose to apply, we speculate that with better communication on what it takes to be a TA and more encouragement to apply, instructors could draw from a wider pool of suitable applicants.

Despite the potential usefulness of this data in hiring, we do not wish to suggest that the TAs who have been hired in the past were unsatisfactory. To the contrary, we have had very effective TA cohorts in the past. Students may engage with the course in different ways – not just through the discussion board – and students who do not engage through discussion board participation can certainly perform TA duties well. As a result, hiring processes should not rely entirely on discussion board engagement data. Other sources of information, such as the candidate’s engagement in related activities or performance in an interview, are also vital. Nevertheless, we argue that there is still value to instructors to have this information, both to enhance the limited information available and to complement the pool of applicants who may have demonstrated strength in other ways. In our own personal experience, it has been difficult, especially with the high rate of enrollment growth in first year courses, to find enough qualified candidates, so broadening the pool is beneficial.

Finally, we should acknowledge an ethical issue that would be raised if we were to use data for hiring without the explicit consent of the people generating it. The discussion board was positioned as an additional resource for students to use for the purpose of learning. They did not consent to have their discussion board activities collected to contribute to a future hiring decision – for a job opportunity that they probably did not know exist. We wish to emphasize that this data was used for research only; the instructors made their hiring decisions before being given additional data to consider. If this data were to be used in hiring in the future, we feel it is necessary to let students know how their data might be used, and that might influence how they behave on the discussion board.

However, once students are aware that their actions on the board may have long-term consequences, they may change how they use the discussion board. Some students, hoping

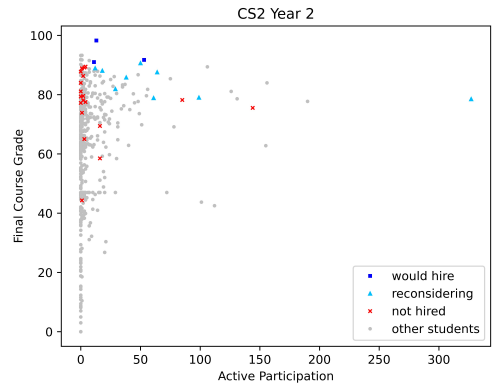
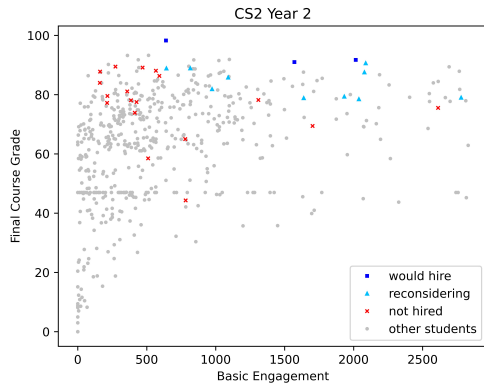
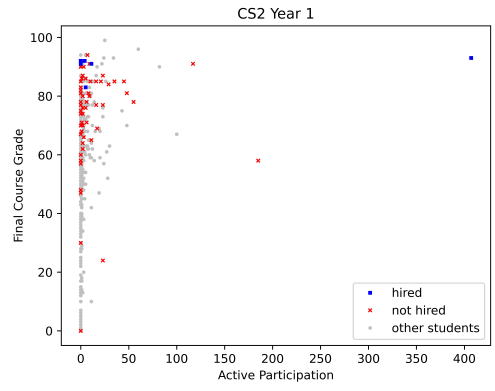
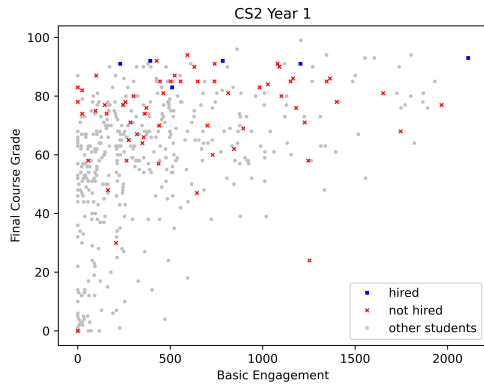
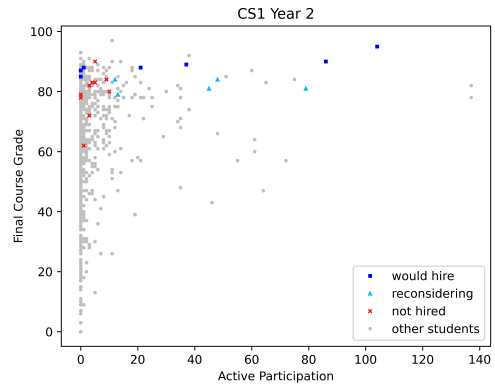
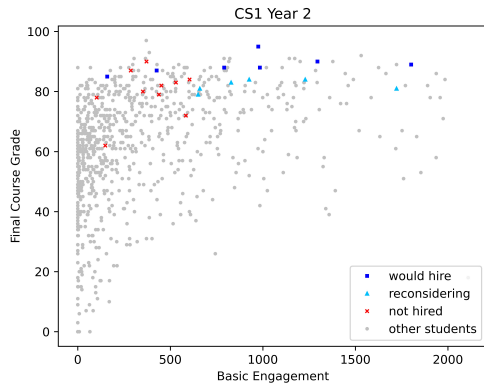
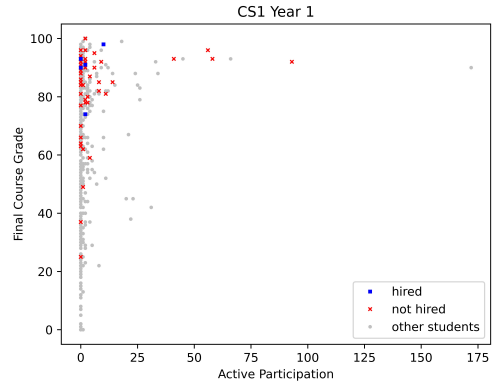
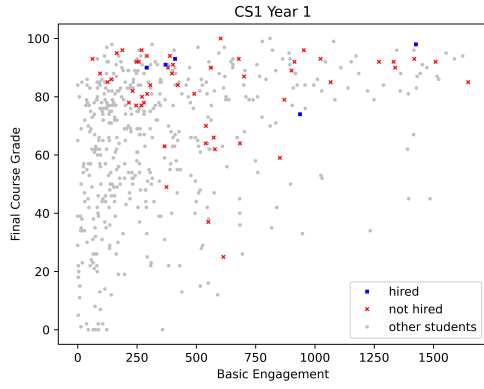


Figure 1. Grades vs. Basic Engagement

Figure 2. Grades vs. Active Participation

to be noticed, may increase their engagement on the board. Overall, we believe this will have a positive impact on the community. While there is a possibility that some of the additional engagement may be superficial, being more engaged in discussions will likely be beneficial, and additional (or quicker) responses would improve the experience for other students. Other students may choose not to engage with the board at all, as they may no longer feel that they benefit from anonymity. However, in the courses we studied, the students could be anonymous to each other but were definitely aware that they were not anonymous to instructors or TAs; as a result, we think that any perceived loss of anonymity will be small. Also, if we publish the metrics we are considering – sustained engagement on the board by reading and posting – and make it a teaching moment about the uses to which the data they generate may be used (e.g., on social media), we believe that students will benefit and will not feel threatened by our particular use of the data. Nevertheless, this is an issue that will require careful consideration during implementation and will require additional research to account for the changes in student behaviour, such as increases in low-quality posts, that are caused by the knowledge that their behaviour is being monitored. In particular, this may require methods for identifying the importance of a post or its quality [1].

CONCLUSIONS

Teaching assistants are important members of the instructional team in large courses and are responsible for much of the mentoring and feedback provided to students. As a result, TAs must be carefully selected, and we believe they should be people who have demonstrated a willingness to engage in the classroom environment in positive ways. In this paper, we have shown that discussion board data can be used to supplement the existing hiring process, particularly in situations where the instructor is unfamiliar with the students in the pool or where the TA candidate pool is smaller than desired. When given discussion board data related to engagement from a previous semester, the instructors we consulted agreed that the information could have been useful and would have led them to at least interview a wider pool of applicants.

In the future, we are considering combining this data with other data available from the course, such as programming exercise submissions, to provide more rich profiles about applicants. We will also explore methods for providing these profiles to students during the course as a teaching tool, so that they are aware of how data they generate might be used, both in the narrow context of TA hiring and more broadly in the context of their use of social media and other “free” online tools. These profiles may also be useful to help students see the kinds of behaviours their instructors believe are correlated with high performance.

REFERENCES

- [1] Noah Bilgrien, Roy Finkelberg, Chirag Tailor, India Irish, Girish Murali, Abhishek Mangal, Niklas Gustafsson, Sumedha Raman, Thad Starner, and Rosa Arriaga. 2019. PARQR: Augmenting the Piazza Online Forum to Better Support Degree Seeking Online Masters Students. In *Proceedings of the Sixth (2019) ACM Conference on Learning@ Scale*. 1–4.
- [2] Urvashi Desai, Vijayalakshmi Ramasamy, and James D Kiper. 2020. A Study on Student Performance Evaluation using Discussion Board Networks. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*. 500–506.
- [3] Ruth B Ekstrom. 1976. Teacher aptitudes, knowledge, attitudes, and cognitive style as predictors of teaching behavior. *Journal of Teacher Education* 27, 4 (1976), 329–331.
- [4] Carrie Demmans Epp, Krystle Phirangee, and Jim Hewitt. 2017. Student actions and community in online courses: The roles played by course length and facilitation method. *Online Learning Journal* 21, 4 (2017).
- [5] David Huffaker. 2010. Dimensions of leadership and social influence in online communities. *Human Communication Research* 36, 4 (2010), 593–617.
- [6] Greg Kearsley and B Schneiderman. 1999. Engagement theory: A framework for technology-based learning and teaching. *Educational Technology* 38, 5 (1999), 20–23.
- [7] Irena Koprinska, Joshua Stretton, and Kalina Yacef. 2015. Students at Risk: Detection and Remediation. In *EDM*. 512–515.
- [8] Jean Lave, Etienne Wenger, and others. 1991. *Situated learning: Legitimate peripheral participation*. Cambridge University Press.
- [9] Murat Oztok, Daniel Zingaro, Alexandra Makos, Clare Brett, and Jim Hewitt. 2015. Capitalizing on social presence: The relationship between social capital and social presence. *The Internet and Higher Education* 26 (2015), 19–24.
- [10] Piazza. 2020. Piazza homepage. (2020). Retrieved July 14, 2020 from <https://piazza.com>
- [11] Anthony G Picciano. 2002. Beyond student perceptions: Issues of interaction, presence, and performance in an online course. *Journal of Asynchronous learning networks* 6, 1 (2002), 21–40.
- [12] Jennifer Preece and Ben Shneiderman. 2009. The reader-to-leader framework: Motivating technology-mediated social participation. *AIS transactions on human-computer interaction* 1, 1 (2009), 13–32.
- [13] Jim Waters. 2012. Thought-Leaders in Asynchronous Online Learning Environments. *Journal of Asynchronous Learning Networks* 16, 1 (2012), 19–34.
- [14] Mengxiao Zhu, Yoav Bergner, Yan Zhang, Ryan Baker, Yuan Wang, and Luc Paquette. 2016. Longitudinal engagement, performance, and social connectivity: a MOOC case study using exponential random graph models. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge*. 223–230.