

Profiling Student Groups in Online Discussion with Network Analysis

Jeon-Hyung Kang, Jihie Kim and Erin Shaw

University of Southern California Information Sciences Institute
4676 Admiralty Way, Marina del Rey, CA, U.S.A
{jeonhyuk, jihie, shaw}@isi.edu

Abstract

As online discussion boards become a popular medium for collaborative problem solving, we would like to understand patterns of group interactions that lead to collaborative learning and better performance. In this paper, we present an approach for assessing collaboration in online discussion, by profiling student-group participation. We use a modularity function to compute optimal discussion group partitions and then examine usage patterns with respect to high-versus low-participating students, and high- versus low-performing students as measured by grades. We apply the profiling technique to a discussion board of an undergraduate computer science course with three semesters of discussion data, comprising 142 users and 1620 messages. Several patterns are identified, and in particular, we show that high achievers tend to act as ‘bridges’, engaging in more diverse discussions with a wider group of peers.

Keywords: Student online discussions, group detection in discussions

Introduction

Online discussion boards play an important role in distance education and web enhanced courses. Studies have shown online discussion to be a promising strategy for promoting collaborative problem solving and discovery-oriented activities, however, student participation can vary highly; some students post only one or two messages during the whole semester, while others participate more often and interact with many other students. Some students communicate with only a limited number of peers, while others interact with a wider group of students and participate in more varied discussion topics. It is difficult to understand the many different types of group interactions that occur in online discussions and even more difficult to understand how they affect collaborative learning.

We would like to identify and understand patterns of group interactions that lead to collaborative learning and better performance. The patterns might be used to develop pedagogical strategies for promoting more desirable interactions and increasing student learning. Most of the existing computational work on qualitative discussion analysis has focused on analyzing dialogue patterns in individual discussion threads (McLaren et al., 2008; Scheuer O.&

McLaren, 2008, Ravi & Kim 2007; Cakir et al., 2005; Feng, Kim, Shaw & Hovy, 2005) or analyzing the impact of tutors in student discussions (Light et al., 2000; Shaw 2005). Although these results provide good hints about student behavior within online discussions, they have not yet yielded insightful information about collaborative learning, such as how groups form during online interactions over multiple discussion threads, and how group interactions affect student learning.

In this paper we present our first steps towards analyzing group activities in online discussions using community modeling algorithms. For each discussion forum, we identify a set of communities based on connections between students and the threads they contribute to. We compare high-participating and low-participating students with respect to their participation patterns within their groups or across different groups. High performing and low-performing students are determined based their project grades. We found that high-participating students who participate in a greater number of discussions don’t necessarily get better grades. However, we see that high performing students tend to act as bridges among several different community groups. That is, their participation is not limited to discussions by one or two groups; they appear to engage in more diverse discussions with a wider group of peers.

Our work takes place in the context of an undergraduate course discussion board that is an integral component of an Operating Systems course in the Computer Science department at the University of Southern California. The course is offered every semester, and always taught by the same lecturer. Each semester, there are four projects based assignments, which remain consistent across semesters, and a final exam. Students use a discussion board, most commonly, to seek help on the project assignments. There is a discussion forum for each of the four projects every semester. There are also administrative, lecture and humor forums. The discussion board is an extended version of phpBB, an open source bulletin board application. Discussion data and project grades from the 2008 spring, summer and fall semesters were used for the presented work.

The paper first present our framework for modeling groups in online discussions, and then presents the re-

sults of analyzing three semesters of data and relating student participation to grades. We report on typical behavior patterns found.

The main contributions of the paper are:

1. Presenting a novel approach for assessing collaboration in online discussion, based on student-group participation profiling.
2. Analyzing how high- and low-performing students from computer science courses interact in groups, and identifies several behavior patterns.

Modeling Discussion Community

The discussion community is modeled based on information about its students and the discussion threads in which they participate. The relationship can be represented as a directed graph, where nodes are either users or threads, and edges connect users who participate in threads (Figure 1). Thus, a discussion graph has $T+U$ nodes and E edges, where T is the number of discussion threads and U is the number of discussants (users), and E represents user participation in a thread. For example, if user i participates in thread j , we generate edge e_{ij} from node i to j (Figure 1).



Figure 1. Users u_1 and u_2 contribute to thread t_1 , along edges $e_{1,1}$ and $e_{2,1}$ in a discussion graph.

The graph is then partitioned to detect optimal communities of discussants. For example, if two students are connected together through many threads, we say they belong to the same group, or community. This is done by first representing the graph as a discussion (student-thread) matrix, and then finding a partitioning that maximizes the strength of the connections within a group, called the *modularity* (Newman, 2006). Once an ideal partitioning is found, we study patterns between and within groups.




Figure 2. Distribution of number of threads per student discussants (Y axis) in the dataset.

To profile discussion participation, we first identify *high*- and *low*-participating students. High-participating students are defined as those who participate in many discussion threads, where low-participating students post only small numbers of messages in one or two threads. Thus, *high* and *low* are with respect to threads as well as messages. To set

the threshold for high-participating vs. low-participating students, we analyzed the degree of participation by individual students in our largest dataset, in which 50 discussants participated in 100 discussion threads (Figure 2). Since 80% of the messages were written by 20% of the students, we labeled the top 20% students high-participating, and the rest low-participating. In the dataset for this work, the top 20% students participate in at least five discussion threads. We applied the power-law to set the threshold.

Matrix Construction

Within the discussion graph, nodes represent both users and threads, with the first n nodes being users and next k being threads, and an edge $e=(i,j)$ denotes user i contributing to thread j , where contributing to (or participating in) means posting a message to a thread for the purpose of either initiating a thread or posting a reply.

An adjacency matrix A is then constructed from the discussion graph, where entry i,j is 1 if and only if there is an edge from node i to node j , and 0 otherwise. If we sum the column values for row i , we get the number of threads in which student i participated. Similarly, the number of users participating in thread j can be calculated by adding up the row values for j column. Sub-matrix TU is the transpose of UT . The values of user i to user i in sub-matrix UU and thread j to thread j in sub-matrix TT are all zero. So all diagonal values of matrix A are zero.

$$A = \begin{bmatrix} UU_{u \times u} & UT_{u \times t} \\ TU_{t \times u} & TT_{t \times t} \end{bmatrix}$$

In this paper, we consider only the relationship from threads to users and users to threads, so sub-matrices UU and TT are zero matrices, however in the future, we plan to include additional group information such as project partners in the sub-matrix UU , and thread relation information such as threads with related topics in the sub-matrix TT .

Community Detection

We apply Newman's (2006) concepts of modularity to generate student groups. The modularity (Q) measures how good the given partition is.

$$Q = (\text{connectivity of community})$$

$$- (\text{expected connectivity of community}).$$

where *connectivity of community* represents the number of edges that lie within groups in the given community, and *expected connectivity of community* is the same quantity when the edges are placed randomly and the vertices have the same degrees. We wish to find community which max-

imizes Q, because positive (higher) modularity values indicate the existence of the community.

The modularity is represented by

$$Q = \frac{1}{2m} \sum_{ij} [A_{ij} - P_{ij}] \delta(g_i, g_j)$$

where m is the total number of edges, A is our adjacency matrix from the previous section, P is the expected connectivity matrix, which is the expected number of edges in an equivalent network where k_i and k_j are degrees of the nodes.

$$P_{ij} = \frac{k_i k_j}{2m}$$

and g_i is the group that i belongs to. If g_i and g_j belong to the same group, $\delta(g_i, g_j) = 1$, otherwise $\delta(g_i, g_j) = 0$. The problem of maximizing Q is NP-Complete, so a heuristic algorithm from Girvan and Newman (2002) was used. Thus, we transformed the problem to an Eigen value problem, and then iteratively divided existing nodes into two groups, while there was a valid division.

Preliminary Results

'Project 2' discussion communities were modeled as a first test. These communities comprised all participating users and threads from three semesters of Project 2 discussion forums. Project 2 was chosen because its forums had the highest number of participants. Descriptive statistics for the Project 2 communities are shown in Table 1. In Table 2, we see that 47-48% of total discussion board users participated in Project 2 discussions in each semester.

semester	The number of forums	The number of Users	The number of Messages	The number of Threads
2008 Spring	9	36	366	122
2008 Summer	9	27	240	77
2008 Fall	9	79	1014	275

Table 1: Discussion data from three courses

semester	The number of Users	The number of Messages	The number of Threads
2008 Spring	17(47%)	119(32%)	25(20%)
2008 Summer	13(48%)	82(34%)	30(26%)
2008 Fall	37(47%)	472(46%)	98(36%)

Table 2: Project 2 forum data only

Table 3 shows the average grades over three semesters for students who participated in discussions compared to those who didn't, during the three semesters. Participating students have higher average grades and the averages are consistent across different semesters. The differences in

the degree of participation or the number of users do not appear to affect scores.

semester	Students using discussion board	Students not using discussion board
2008 Spring	34.45/40.00	32.07/40.00
2008 Summer	34.46/40.00	32.20/40.00
2008 Fall	35.20/40.00	33.00/40.00

Table 3: Average grades: discussion board users vs. non users

Groups Identified in Discussions

Figures 3-5 show groups found in the discussions from the three courses. The light blue boxes compass the different groups, ellipses represent students and small yellow rectangles represent threads. Students with lower grades have red ellipses and students with higher grades have blue ellipses. The fall 2008 discussions show more groups and heavier interactions among students.

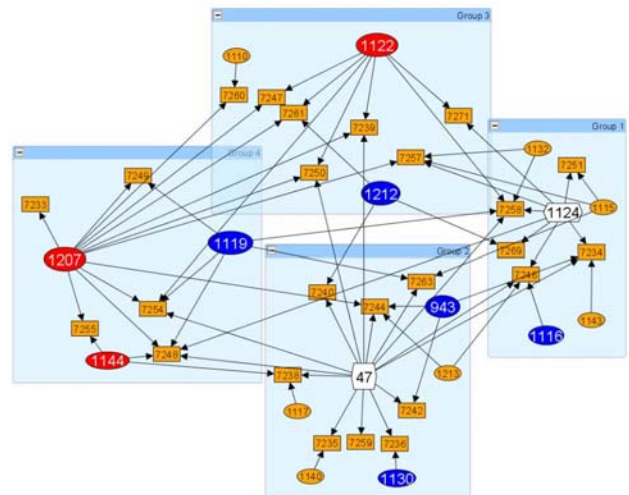


Figure 3: Group distribution in spring 2008 semester.

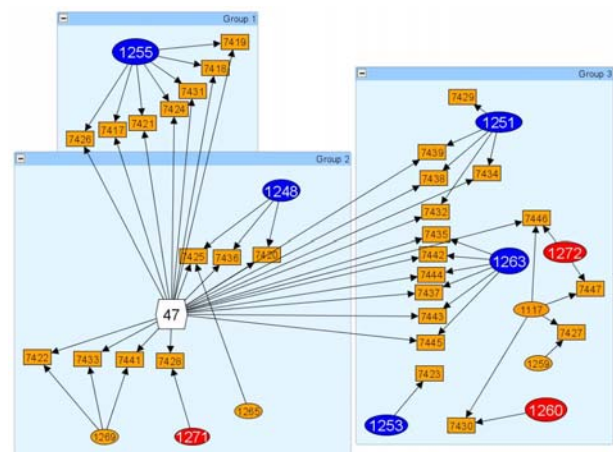


Figure 4: Group distribution in summer 2008 semester.

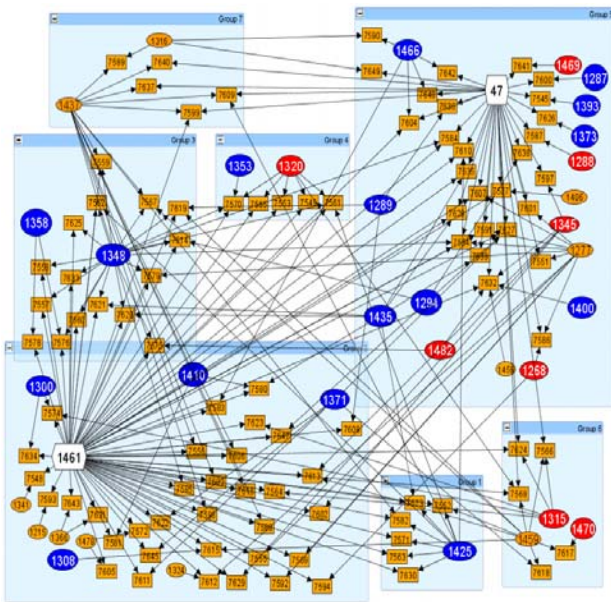


Figure 5: Group distribution in fall 2008 semester.

Profiling Students with Identified Groups

	High Participating Student Students	Low Participating Student Students	Active Group Participants (AGP)	
Average Score	34.8	34.8	33.6	
	Spring 2008	Summer 2008	Fall 2008	Year 2008
High Participating Student Average	28.3	37	35.8	34.8
Low Participating Student Average	36.1	33.3	34.8	34.8

Table 4: High- versus low- participating students.

We compared the grades of high- and low-participating students. High-participating students were defined as those who contributed to at least five threads as described above. As shown in Table 4, there was no difference in the average grades of high- and low-participating students.

Active group participants (AGP) are students who participated in at least five different discussion threads in any given group. The AGP in our data received lower average grades than either high- or low-participating students. This is consistent with our earlier findings that sometimes high participating students are help seekers and not necessarily high performers. A full understanding will require an in-depth investigation of thread features, taking into account the technical quality of messages, dialogue patterns, technical terms used, and discussion topics. We describe our study of individual student interactions below.

	Spring 2008	Summer 2008	Fall 2008	2008 total
Bridge Students	37.3	-	39.3	38.4
Other students	33.4	34.2	34.6	34.3
	Bridge students		Other students	
Average Score	38.4		34.3	

Table 5: Bridge students vs. others

Some students participated in multiple discussion threads across several different groups. We call these students *bridge students*. A bridge student participates in more than three threads which belong to more than three different groups. We have three bridge students in spring 2008 (students 1119, 1212, and 943) and four students in fall 2008 semester (students 1289, 1294, 1435, and 1371). These students are higher-performing students than their peers. The average score of our bridge students is 38.4 out of 40.

Additional Group Interaction Patterns

Spring 2008 Semester

Two AGP of Group 1 and Group 2 are the instructor and the teaching assistance; another two AGP are student 1207 and student 1122. We found that these two students interact with each other across many threads, and both received poor grades. In particular, student 1122 tends to ask about technical concepts taught in the class or ask questions about other students' messages. This could indicate that he/she didn't have a good background in the project area or didn't have the language skills to understand messages posted by other students. Student 1207 also participated in many threads by posting replies or answers but more often he/she just provided wrong answers. For some replies, the instructor sent messages that corrected his/her answers. In contrast to these two students, student number 1119, 1212, and 943 received good grades and participated in several different discussion groups as a bridge. For example, student 1212 participated in three threads across three different groups.

Spring 2008 Semester

When we compare students' behaviors and patterns of summer 2008 to spring 2008, the average number of threads per students in summer 2008 is higher than that in spring 2008 semester. In this semester, often a question and an answer were paired a student asks a question and the instructor provides an answer. There were only two threads where more than three users participated. The instructor's frequent replies may have limited opportunities for other students to answer. Unlike in other semesters, no one received a perfect score (40/40). Finally, there were no bridge students this semester.

Fall 2008 Semester

There were seven groups for the project 2 discussions in the fall 2008 semester. Each group has more than one AGP. Students 1289, 1435, and 1371 were bridge students and all of them received the perfect project assignment score (40/40). Although students 1425 and 1320 were AGP students for the group, only student 1425 received a good grade. Interestingly, these two students contributed to discussions in very different ways. In five out of six threads, student 1425 provided an answer peer queries. In contrast, student 1320 questions were all about basic programming details or easy code debugging problems. Even when students participate in discussions with similar quantitative characteristics, the qualitative characteristics of the message content is necessary to differentiate student intentions. In future work, we plan to include technical quality of messages, dialogue patterns in threads, and discussion topics in the analysis.

Related Work

There has been much prior work on tutorial dialogue (e.g. Graesser et al., 2001, Tetreault et al., 2008). Although some of the techniques are closely related, most of them focus on spoken dialogue or conversation in tutoring systems rather than threaded discussions. Recently, some researchers worked on qualitative assessment of discussions including student reasoning (McLaren et al., 2007). Our work provides complementary capability by providing a group-based interaction assessment framework.

There has been increasing interest in online dialogue including email message analysis (Lampert et al. 2008) and effort to analyze interactions in on-line communities. For example, Talk-to-me (Arguello et al., 2006) can predict the likelihood that a message will receive a reply based on the content of the message and the message sender. However, student discussions tend to focus on problem solving rather than task request and commitment common in email message threads in project management applications.

There have been significant advances in understanding networks using community detection algorithms. Detecting communities is based on the idea that if there are more edges among group of vertices than those from remainders of network, that group of vertices can be a community. Recently, community modeling techniques have been applied to email messages (Diehl et al., 2008) and discussion. For example, Park and Maurer (2009) applied clustering to identify consensus and consensus facilitators. Although the techniques may be used to determine if student discussions reach a conclusion, in assessing student activities, we need to additionally incorporate inter-group activities, as our results indicate. Topic, joint group and role discovery in social networks were also presented by McCallum et al. (2005) using Latent Dirichlet Allocation

however their work focuses on sender and receiver roles only and does not use community modeling techniques.

Discussion and Future Work

For modeling social structure in discussion board discussions, community detection algorithms are not sufficient. Clustering algorithms may need to be merged to identify similarity between threads and users. Most of the work was based on un-weighted node to node connections; however some discussion topics are more important than others, and users might be weighted based on their roles or expertise (e.g., instructors, mentors). Additionally, we might weight threads based on the number of technical-terms, the number of replies or the quality of a discussion. We plan to incorporate these factors to strengthen our community detection algorithm. On discussion boards, user-to-user relations and thread-to-thread relations may be as important as user-to-thread relations. By adding more edges between users and threads, we expect to find more interesting community groups.

Moreover, there are message-level features that we can potentially incorporate in our model; for example, speech acts represent roles that individual messages play (questions, answers, objections, elaborations, etc.) (Ravi and Kim 2007). Students who send questions can be represented differently than students who send elaborate answers. The technical quality of the messages also seems to be an important factor in deciding the role of the student.

Acknowledgements

The authors are indebted to Senior Lecturer Dr. Michael Crowley for his continuous support of the project. The work was funded under a National Science Foundation Course, Curriculum, and Lab Instruction Phase 2 grant (#0618859).

References

- [1] Arguello, J., Butler, B. S., Joyce, L., Kraut, R., Ling, K. S., & Wang, X. (In press) Talk to me: Foundations for successful individual-group interactions in online communities. *In CHI 2006: Proceedings of the ACM conference on human-factors in computing systems*.
- [2] Cakir, M., Xhafa, F., Zhou, N., and Stahl, G. Thread-based analysis of patterns of collaborative interaction in chat, *AI in Education Conference*, 2005.
- [3] Kim, J. and Shaw, J.: Pedagogical Discourse: Connecting Students to Past Discussions and Peer Mentors within an Online Discussion Board, *The 21st Innovative Applications of Artificial Intelligence Conference (IAAI-2009)*, 2009.

- [4] Ravi, S., Kim, J.: Profiling Student Interactions in Threaded Discussions with Speech Act Classifiers, *Proceedings of AIED-2007*.
- [5] Shaw E., Kim J. and Suparnakoon P.: Mentor Match: Using student mentors to scaffold participation and learning within an online discussion board, *Proceedings of AIED-2009*.
- [6] Graesser, A., VanLehn, K., Rosé, C., Jordan, P., Harter, D. 2001. Intelligent Tutoring Systems with Conversational Dialogue. *AI Magazine* 22(4).
- [7] Lampert, A., Dale, R., and Paris, C. 2008. The Nature of Requests and Commitments in Email Messages, *AAAI workshop on Enhanced Messaging*.
- [8] McCallum, A. Corrada-Emmanuel, A., Wang, X., 2005. Topic and Role Discovery in Social Networks, *Proceedings of International Joint Conference on AI*.
- [9] McLaren, B., Scheuer, O., De Laat, M., Hever, R., De Groot, R. & Rose, C. 2007. Using Machine Learning Techniques to Analyze and Support Mediation of Student E-Discussions, *Proceedings of AIED-2007*.
- [10] M. Newman.: Finding community structure in networks using the eigenvectors of matrices, *Physical Review E*, vol. 74, p. 036104, 2006.
- [11] Katz, L. 1953. A New Status Index Derived From Sociometric Analysis, *Psychometrika*, vol. 18, pp. 39-40.
- [12] Light, V., Nesbitt, E., Light, P., & Burns, J.R. (2000) 'Let's You and Me Have a Little Discussion': computer mediated communication in support of campus-based university courses. *Studies in Higher Education*, Volume 25, No. 1
- [13] M. E. J. Newman.: Detecting community structure in networks, *Eur. Phys. J. B* 38, 321-330, 2004
- [14] M Girvan and M E. J. Newman, Community structure in social and biological networks, *Proc. Natl. Acad. Sci. USA* 99, 7821-7826, 2002
- [15] Shelly Park and Frank Maurer.: A Network Analysis of Online Forum Discussions on Executable Acceptance Test Driven Development, 2009
- [16] Pallof, R.M., Pratt, K. 1999. Building Learning Communities in Cyberspace: Effective Strategies for the online classroom
- [17] Feng, D., Kim, J., Shaw, E. (2005). Towards Modeling Threaded Discussion using Induced Ontology Knowledge, *Proceedings of the National Conference on Artificial Intelligence (AAAI-2006)*.
- [18] Scheuer O. and McLaren B.M., Helping teachers handle the flood of data in online student discussions, *In proceedings of 9th International Conference on Intelligent Tutoring Systems, 2008 Montreal, Canada. Lecture Notes in Computer Science, Vol. 5091, pp.323-332, 2008.*