

Transactivity as a Predictor of Future Collaborative Knowledge Integration in Team-Based Learning in Online Courses

Miaomiao Wen, Keith Maki, Xu Wang, Steven P Dow, James Herbsleb and Carolyn Rose
Carnegie Mellon University
Pittsburgh, USA
{mwen,kmaki,xuwang,spdown,jdh,cprose}@cs.cmu.edu

ABSTRACT

To create a satisfying social learning experience, an emerging challenge in educational data mining is to automatically assign students into effective learning teams. In this paper, we utilize discourse data mining as the foundation for an online team-formation procedure. The procedure features a deliberation process prior to team assignment, where participants hold discussions both to prepare for the collaboration task and provide indicators that are then used during automated team assignment. We automatically assign teams in a way that maximizes average observed pairwise transactivity exchange within teams, whereas in a control condition, teams are assigned randomly. We validate our team-formation procedure in a crowdsourced online environment that enables effective isolation of variables, namely Amazon's Mechanical Turk. We compare group knowledge integration outcomes between the two team assignment conditions. Our results demonstrate that transactivity-based team assignment is associated with significantly greater knowledge integration ($p < .05$, effect size 3 standard deviations).

1. INTRODUCTION

Although there are typically thousands of students in a Massive Open Online Course (MOOC), social isolation is still the norm in the current generation of MOOCs. However, there is evidence that many students would prefer to have more social engagement in that context. Recent research shows that a quarter of learners want to meet new people in their courses; and another 20% of learners in typical MOOCs want to take their courses with friends or colleagues [17]. To satisfy learners' social needs, there is growing interest in enabling group learning in MOOC learning contexts. Recent emerging platforms like NovoEd¹ and cMOOCs are designed with team-based learning or social interaction at center stage. Additionally, many recent xMOOCs are adopting

¹<https://novoed.com>

team-based learning features (e.g., in EdX²). There is accumulating evidence that social interaction is associated with enhanced commitment to the course [11], which has the potential to address one of MOOC critics' biggest concerns, namely high attrition rates [18]. However, how to automatically assign students to effective MOOC learning groups is still an open question [12, 25, 20]. Methods for mining educational data have been used to optimize instruction or feedback for individuals [21]. In this paper we explore how a form of educational data mining (namely, mining of discussion behavior) can be used to optimize the experience of collaborative learners through the support of effective team formation.

Algorithms for group assignment typically bring together students based on learning style, personality or demographic information. For team assignments based on such algorithms, student information must be collected and then provided to the algorithm [9]. Because of the paucity of available student personal information in MOOCs, designing a team-formation process that relies on mining of discussion data to fill in missing information would be a valuable contribution. Moreover, research identifying valuable evidence for effective team formation is needed since recent work shows that forming teams based on typical demographic features, e.g. gender and time zone, does not significantly improve teams' engagement and success in MOOCs [25]. In an online interaction, demographic information about learners is only relevant to the extent that it influences how those students come across and interact with others. Thus, observation of behavior and interaction between students may be a better source of insight for assigning students to groups in which they will function well as a team. This provides an excellent opportunity for data mining technology to make a contribution in support of valued learning processes. The alternative to automated assignment is self-selected teams. When a student population is large, which is usually the case in MOOCs, it is difficult for students to navigate through a list of students or teams to find a team that fits. Previous work has shown that many self-selected teams fail in team-based MOOCs [23]. As an alternative to both of these approaches, we design a practical group-formation procedure through which participants are organized into small groups

²https://courses.edx.org/courses/course-v1:McGillX+GROOCx+T3_2015,
<https://www.edx.org/course/medicinal-chemistry-molecular-basis-drug-davidsonx-d001x-1>

based on the data mined from their participation processes in the course. This procedure uses a deliberation process, where participants hold discussions in preparation for the collaboration task; teams are then automatically assigned based on features of their interaction during deliberation.

In recent years there has been increasing interest in mining discourse data for insights into learning processes [7], for understanding factors associated with attrition in MOOCs [16], and for building models to trigger dynamic support for collaborative learning [11]. In this paper, we mine students' collaborative process to collect information for automatic team assignment. In particular, we automatically identify an important property of discourse, transactivity, from students' discussion. Transactivity is known to be higher within groups where there is mutual respect [5] and a desire to build common ground [14]. Previous studies showed that high transactivity groups are associated with higher learning [22], higher knowledge transfer [13], and better problem solving [5]. Prior work has demonstrated success at automatic detection of transactivity and relevant discussion constructs [14]. Because of the social underpinnings of transactivity, it is reasonable to hypothesize that automated detection of transactivity could form the basis for an automated group assignment procedure in online learning contexts. In this paper, we combine text-mining and algorithm-based team formation; We study whether by grouping individuals with a history of engaging in more transactive communication during a pre-collaboration deliberation can help them achieve more effective collaboration in their teams. Simply stated, our research question is:

Can evidence of transactive discussions during deliberation inform the formation of more successful teams?

As a step towards effective team-based learning in MOOCs, in this paper, we explore the team-formation process in an experimental study conducted in an online setting that enables effective isolation of variables, namely Amazon's Mechanical Turk (MTurk). While crowd workers likely have different motivations from MOOC students, their remote individual work setting without peer contact resembles today's MOOC setting where most students learn in isolation [6]. This allows us to test the causal connection between variables in order to identify principles that later we will test in an actual MOOC. A similar approach was taken in prior work to inform design of MOOC interventions for online group learning [6]. We designed a collaborative knowledge integration task where participants work together on writing an energy proposal for a city. This knowledge integration task is modeled after ones used in earlier collaborative learning studies [4]. The participants in our study will be referred to as students throughout the paper.

2. METHODS

Our experimental study is designed as a validation of a team-formation paradigm. In this paradigm, we attempt to offer teams a running start in their collaboration work by starting them with individual work, which they then discuss as a community. In addition to providing the basis for assignment to teams, the community engagement prior to team formation provides students with a breadth of exposure to different perspectives relevant to the group work. Based

on the interactions displayed during this community discussion, students are automatically assigned to teams. The students then enter their teams for the bulk of their group work. We test a transactivity-maximization team-formation method. Instead of grouping students high in transactivity into teams and students low in transactivity together, the team assignment algorithm maximizes the average amount of transactive communication within all the teams through a constraint satisfaction algorithm.

2.1 Experimental Paradigm

2.1.1 Collaboration Task Description

For the team task, we designed a highly-interdependent collaboration task that requires negotiation in order to create a context in which effective group collaboration would be necessary for task success. The task is comparable to a course project where a student team writes a proposal collaboratively. We used a Jigsaw paradigm, which has been demonstrated as an effective way to achieve a positive group composition and is associated with positive group outcomes [4]. In a Jigsaw task, each student is given a portion of the knowledge or resources needed to solve the problem, but no one has enough to complete the task alone. Following the Jigsaw paradigm, each member of the team was given special knowledge of one of the four energy sources, and was instructed to represent the values associated with their energy source in contrast to the rest, e.g. coal energy was paired with an economical energy perspective. The team collaborative task was to select a single energy plan and write a proposal arguing in favor of the group decision with respect to the associated trade-offs, meaning team members needed to negotiate a prioritization among the city requirements with respect to the advantages and disadvantages they were cumulatively aware of. The set of potential energy plans was constructed to reflect different trade-offs among the requirements, with no plan satisfying all of them perfectly. This ambiguity created an opportunity for intensive exchange of perspectives. The collaboration task is shown in Figure 1.

2.1.2 Experimental Procedure

We designed a four-step process for the task:

Step 1: Preparation. In this step, each student was asked to provide a nickname, which would be used in the deliberation and collaboration phases. To prepare for the Jigsaw task, each student was randomly assigned to read an instructional article about the pros and cons of a single energy source. Each article was approximately 500 words, and covered one of four energy sources (coal, wind, nuclear, and hydro power). To strengthen their learning and prepare them for the proposal writing, we asked them to complete a quiz reinforcing the content of their assigned article. The quiz consisted of 8 single-choice questions, and feedback including correct answers and explanations was provided along with the quiz.

Step 2: Pre-task. In this step, we asked each student to write a proposal to recommend one of the four energy sources (coal, wind, nuclear, and hydro power) for a city given five requirements, e.g. "The city prefers a stable energy". After each student finished this step, their proposal was automatically posted in a forum as the start of a thread with the title "[Nickname]'s Proposal".

In this final step, you will work together with other Turkers to recommend a way of distributing resources across energy types for the administration of City B. City B requires 12,000,000 MWh electricity a year from four types of energy sources: coal power, wind power, nuclear power and hydro power. We have provided 4 different plans to choose from, each of which emphasizes one energy source as primary. Your team needs to negotiate which plan is the best way of meeting your assigned goals, given the city’s requirements and information below.

City B’s requirements and information:

1. City B has a tight yearly energy budget of \$900,000K. Coal power costs \$40/MWh. Nuclear power costs \$100/MWh. Wind power costs \$70/MWh. Hydro power costs \$100/MWh.
2. The city is concerned with chemical waste. If the main energy source releases toxic chemical waste, there is a waste disposal cost of \$2/MWh.
3. The city is a famous tourist city for its natural bird and fish habitats.
4. The city is trying to reduce greenhouse gas emissions. If the main energy source releases greenhouse gases, there will be a “Carbon tax” of \$10/MWh of electricity.
5. The city has several large hospitals that need a stable and reliable energy source.
6. The city prefers renewable energy. If renewable energies generate more than 30% of the electricity, there will be a renewable tax credit of \$1/MWh for the electricity that is generated by renewable energies.
7. The city prefers energy sources whose cost is stable.
8. The city is concerned with water pollution.

	Energy Plan				Cost	Waste disposal cost	Carbon tax	Renewable tax credit	Total
	Coal	Wind	Nuclear	Hydro					
Plan 1	40%	20%	20%	20%	\$840,000K	\$14,400K	\$48,000K	\$9,600K	\$892,800K
Plan 2	20%	40%	20%	20%	\$912,000K	\$0	\$0	\$11,000K	\$901,000K
Plan 3	20%	20%	40%	20%	\$984,000K	\$14,400K	\$0	\$9,600K	\$988,800K
Plan 4	20%	20%	20%	40%	\$984,000K	\$0	\$0	\$11,000K	\$973,600K

Figure 1: This figure displays the collaborative task as it was presented to the students. In addition to the task statement, they had a chat interface and a shared document space to work in.

Step 3: Deliberation. In this step, students joined a threaded forum discussion akin to those available in many online environments. Each proposal written by the students in the Pre-task (Step 2) was displayed for students to read and comment on. Each student was required to write at least five replies to the proposals posted by the other students. To encourage the students to discuss transactively, the task instruction for this step included the request to, when replying to a post, “elaborate, build upon, question or argue against the ideas presented in that post, drawing from the argumentation in your own proposal where appropriate.”

Step 4: Collaboration. In the collaboration step, team members in a group were first gathered for synchronous interaction and then directed to a shared document space to write a proposal together to recommend one of four suggested energy plans based on a city’s eight requirements. Students in the same team were able to see each other’s edits in real time, and were able to communicate with each other using a synchronous chat utility on the right sidebar. The collaborative task was designed to contain richer information than the individual proposal writing task in Step 2.

2.1.3 Outcome Measures

We evaluated team success using two types of outcomes, namely objective success through quantitative task performance (i.e., the quality of the integrated proposal, which indicates collaborative knowledge integration [3]) and process measures, as well as subjective success through a group satisfaction survey. The quantitative task performance measure was an evaluation of the quality of the proposal produced by the team. The goal of evaluating the team knowledge integration process is to distinguish instances when students are

making statements based on reasoning from simply repeating what they have read. In particular, the scoring rubric defined how to identify the following elements for a proposal: (1) Which requirements were considered; (2) Which comparisons or trade-offs were made; (3) Which additional valid desiderata were considered beyond stated requirements; (4) Which incorrect statements were made about requirements. Positive points were awarded to each proposal for correct requirements considered, comparisons made, and additional valid desiderata. Negative points were awarded for incorrect statements. We measured *Team Knowledge Integration* by the total points assigned to the team proposal, i.e. team proposal score. Two PhD students who were blind to the conditions applied the rubric to five proposals (a total of 78 sentences) and the inter-rater reliability was good ($Kappa = 0.74$). The two raters then coded all the proposals.

We used the *length of chat discussion* during teamwork as a measure of team process in the Collaboration step. On average the longer discussions referred to more substantive issues.

Group Experience Satisfaction was measured using a four item group experience survey administered to each student after the Collaboration step. The survey was based on items used in prior work [19, 6]. In particular, the survey instrument included items related to:

- Satisfaction with team experience.
- Satisfaction with proposal quality.
- Satisfaction with the group communication.
- Perceived learning through the group experience.

Each of the items was measured on a 7-point Likert scale.

2.1.4 Control Variables

Intuitively, students who display more effort in the Pre-task might perform better in the collaboration task, so that level of effort is an important control variable. We used each student's individual Pre-task proposal length as a control variable for Individual Performance. Analogously, we used each group's average group member Pre-task proposal length as a control variable for the group knowledge integration analyses.

2.1.5 Transactivity Annotation, Prediction, and Measurement

To enable us to use counts of transactive contributions as evidence to inform an automated group assignment procedure, we needed to automatically judge whether a reply post in the Deliberation step was transactive or not using machine learning. A transactive contribution displays the author's reasoning and connects that reasoning to material communicated earlier. Two example posts illustrating the contrast are shown below:

- Transactive
"Nuclear energy, as it is efficient, it is not sustainable. Also, think of the disaster probabilities".
- Non-transactive
"I agree that nuclear power would be the best solution".

Using a validated and reliable coding manual for transactivity from prior work [14], an annotator previously trained to apply that coding manual annotated 426 reply posts collected in pilot studies we conducted in preparation for the studies reported in this paper. Each of those posts was annotated as either "transactive" or "non-transactive". 70% of them were transactive.

Automatic annotation of transactivity has been reported in the Computer Supported Collaborative Learning literature. For example, researchers have applied machine learning using text, such as chat data [15] and transcripts of whole group discussions [2]. We trained a Logistic Regression model with L2 regularization using a set of features consisting of single word features (i.e., unigrams) as well as a feature indicating the post length [10]. We evaluated our classifier with a 10-fold cross validation and achieved an accuracy of 0.843 and a 0.615 Kappa. Given the adequate performance of the model, we used it to predict whether each reply post in the Deliberation step was transactive or not.

To measure the amount of transactive communication between two students in the Deliberation step, we counted the number of times a pair of their posts in a same discussion thread were transactive; or one of them was a thread starter and the other student's reply was transactive.

2.2 Transactivity Maximization Grouping

The Transactivity Maximization teams were formed so that the average amount of transactive discussion observed in the Deliberation step among the team members in the team

was maximized. A Minimal Cost Max Network Flow algorithm was used to perform this constraint satisfaction process [1]. This network flow algorithm tackles resource allocation problems with constraints. In our case, we need to satisfy the Jigsaw constraint. At the same time, the minimal cost part of the algorithm maximized the transactive communication that was observed among the group members during the Deliberation step. The algorithm finds an approximately optimal grouping within $O(N^3)$ (N = number of students) time complexity. A brute force search algorithm, which has an $O(N!)$ time complexity, would take too long to finish in real time.

Our algorithm can achieve an approximately optimal solution in an admissible time. Instead of maximizing the pair-wise accumulated transitivity post count, we approximate the solution by maximizing the accumulated transitivity post count between two adjacent pairs of users. The algorithm can be generalized to form teams of any size. In our experiment, the team size is 4. We build a directed weighted graph based on students' discussion network. Then we use the successive shortest path algorithm to find a sub-optimal, but nevertheless substantially better than random grouping [1]. The algorithm greedily finds a flow with minimum cost until there is no remaining flow in the network, as outlined in Algorithm 1.

Algorithm 1 Successive Shortest Paths for Minimum Cost Max Flow

```
 $f(v_1, v_2) \leftarrow 0 \forall (v_1, v_2) \in E$   
 $E' \leftarrow a(v_1, v_2) \forall (v_1, v_2) \in E$   
while  $\exists \Pi^* \in G' = (V, E')$   
s.t.  $\Pi^*$  a minimum cost path from  $S$  to  $D$  do  
  for each  $(v_1, v_2) \in \Pi^*$   
    if  $f(v_1, v_2) > 0$  then  
       $f(v_1, v_2) \leftarrow 0$   
      remove  $-a(v_2, v_1)$  from  $E'$   
      add  $a(v_1, v_2)$  to  $E'$   
    else  
       $f(v_1, v_2) \leftarrow 1$   
      remove  $a(v_1, v_2)$  from  $E'$   
      add  $-a(v_2, v_1)$  to  $E'$   
    end  
  end  
end
```

2.2.1 Experimental Manipulation

In our study, students participated in a deliberative discussion as a community in a threaded discussion forum prior to being assigned to teams automatically. We investigated how the nature of the experience in that context may contribute to the success of the teams. We made use of a Jigsaw paradigm in the team assignment of teams in both the experimental and control conditions. In the experimental condition, which we termed the Transactivity Maximization condition, we additionally applied a constraint that preferred to maximize the extent to which students assigned to the same team had participated in automatically detected transactive exchanges in the deliberation. In the control condition, which we termed the Random condition, apart from enforcing the Jigsaw constraint, teams were formed by random assignment. In this way we tested the hypothesis that observed transactivity is an indicator of potential for effective

tive team collaboration. We ran the study in 10 separate batches, with 5 batches in each condition. In each batch, all the students in that batch were assigned to teams using the Random strategy or all the students were assigned to teams using the Transactivity Maximization strategy. The average level of amount of transactivity during the deliberation stage was not significantly different between batches. Thus we can test if the team-formation method can predict future collaborative knowledge integration. All the steps and instructions of the task were identical for the two conditions, as described in 2.1.2.

2.3 Participants

Participants were recruited on MTurk with the qualifications of having a 95% acceptance rate on 1,000 tasks or more. Each student was only allowed to participate once. A total of 246 students participated in the experiment, the students who were not assigned into groups or did not complete the group satisfaction survey were excluded from our analysis. The experiment lasted on average 35.9 minutes. We included only teams of 4 students in our analysis. There were in total 27 Transactive Maximization teams and 27 Random teams, with no significant difference in attrition between conditions ($\chi^2(1) = 1.46, p = 0.23$). The dropout rate of students in Random groups was 27%. The dropout rate of students in Transactivity Maximization groups was 19%.

3. RESULTS

As a manipulation check, we compared the average amount of transactivity observed among teammates during the deliberation between the two conditions using a t-test. The groups in the Transactive Maximization condition ($M = 12.85, SD = 1.34$)³ were observed to have had significantly more transactive exchanges during the deliberation than those in the Random condition ($M = 7.00, SD = 1.52$) ($p < 0.01$), with an effect size of 3.85 standard deviations, demonstrating that the maximization was successful in manipulating the average experienced transactive exchange within teams between conditions.

Teams that experienced greater transactivity during deliberation demonstrate better team knowledge integration.

To assess whether the Transactivity Maximization condition resulted in more effective teams, we tested for a difference between group-formation conditions on Team Knowledge Integration. We built an ANOVA model with Grouping Criteria (Random, Transactivity Maximization) as the independent variable and Team Knowledge Integration as the dependent variable. Average team member Pre-task proposal length was again the covariate. There was a significant main effect of Grouping Criteria ($F(1,52) = 6.13, p < 0.05$) on Team Knowledge Integration such that Transactivity Maximization teams ($M = 11.74, SD = 0.67$) demonstrated significantly better performance than the Random groups ($M = 9.37, SD = 0.67$) ($p < 0.05$), with an effect size of 3.54 standard deviations, which is a large effect. Effect size is measured in terms of Cohen's d .

Across the two conditions, observed transactive communication during deliberation was significantly correlated with Team Knowledge Integration ($r = 0.26, p < 0.05$). This

³SD is short for standard deviation in this paper.

also indicated teams that experienced more transactive communication during deliberation demonstrated better Team Knowledge Integration.

Teams that experienced greater transactivity during deliberation demonstrate more intensive interaction within their teams.

In the experiment, students were assigned to teams based on observed transactive communication during the deliberation step. Assuming that individuals that were able to engage in positive collaborative behaviors together during the deliberation would continue to do so once in their teams, we would expect to see evidence of this reflected in their observed team process. Group processes have been demonstrated to be strongly related to group outcomes in face-to-face problem solving settings [24]. Thus, we should consider evidence of a positive effect on group processes as an additional positive outcome of the experimental manipulation.

In order to test whether such an effect occurred, we built an ANOVA model with Grouping Criteria (Random, Transactivity Maximization) as the independent variable and length of chat discussion during teamwork as the dependent variable. There was a significant effect of Grouping Criteria on length of discussion ($F(1,45) = 9.26, p < 0.005$). Random groups ($M = 20.00, SD = 3.58$) demonstrated significantly shorter discussions than Transactive Maximization groups ($M = 34.52, SD = 3.16$), with an effect size of 4.06 standard deviations.

Survey results

For each of the four aspects of the group experience survey, we built an ANOVA model with Grouping Criteria (Random, Transactivity Maximization) as the independent variable and the survey outcome as the dependent variable. Team ID and assigned energy condition (Coal, Wind, Hydro, Nuclear) were included as control variables nested within condition. There were no significant effects on Satisfaction with team experience or with proposal quality. However, there was a significant effect of condition on Satisfaction with communication within the group ($F(1,112) = 4.83, p < 0.05$), such that students in the Random teams ($M = 5.12, SD = 1.7$) rated the communication significantly lower than those in the Transactivity Maximization teams ($M = 5.69, SD = 1.51$), with effect size 0.38 standard deviations. Additionally, there was a marginal effect of condition on Perceived learning ($F(1,112) = 2.72, p = 0.1$), such that students in the Random teams ($M = 5.25, SD = 1.42$) rated the perceived benefit to their understanding they received from the group work lower than students in the Transactivity Maximization teams ($M = 5.55, SD = 1.27$), with effect size 0.21 standard deviations. Thus, with respect to subjective experience, we see advantages for the Transactivity Maximization condition, but the results are weaker than those observed for the objective measures. Nevertheless, these results are consistent with prior work where objectively measured learning benefits are observed in high transactivity teams [8].

4. DISCUSSION

In this paper we presented an experiment to address our research question regarding the extent to which benefit could be achieved by selecting teams based on evidence of trans-

active exchange observed during the deliberation. We designed an automatic team-formation process that combines discourse data mining and algorithm-based team formation. Here we found that teams formed such that observed transactive interactions between team members in the deliberation was maximized displayed objectively better knowledge integration than teams assigned randomly. On subjective measures we see a significant positive impact of transactivity maximization on perceived communication quality and a marginal impact on perceived enhanced understanding, both of which are consistent with what we would expect from the literature on transactivity where high transactivity teams have been demonstrated to produce higher quality outcomes and greater learning [22]. These results provide positive evidence in favor of a design for a team-formation strategy in two stages: Individuals first participate in a pre-teamwork deliberation activity where they explore the space of issues in a context that provides beneficial exposure to a wide range of perspectives. Individuals are then grouped automatically through a transactivity detection and maximization procedure that uses communication patterns arising naturally from community processes to inform group formation with an aim for successful collaboration.

This research was supported in part by funding from Google and the Gates foundation.

5. REFERENCES

- [1] R. K. Ahuja and J. B. Orlin. A fast and simple algorithm for the maximum flow problem. *Operations Research*, 37(5):748–759, 1989.
- [2] H. Ai, R. Kumar, D. Nguyen, A. Nagasunder, and C. P. Rosé. Exploring the effectiveness of social capabilities and goal alignment in computer supported collaborative learning. In *Intelligent Tutoring Systems*, pages 134–143. Springer, 2010.
- [3] M. Alavi and A. Tiwana. Knowledge integration in virtual teams: The potential role of kms. *Journal of the American Society for Information Science and Technology*, 53(12):1029–1037, 2002.
- [4] E. Aronson. *The jigsaw classroom*. Sage, 1978.
- [5] M. Azmitia and R. Montgomery. Friendship, transactive dialogues, and the development of scientific reasoning. *Social development*, 2(3):202–221, 1993.
- [6] D. Coetzee, S. Lim, A. Fox, B. Hartmann, and M. A. Hearst. Structuring interactions for large-scale synchronous peer learning. In *CSCW*, 2015.
- [7] M. Dascalu, S. Trausan-Matu, D. S. McNamara, and P. Dessus. Readerbench: Automated evaluation of collaboration based on cohesion and dialogism. *International Journal of Computer-Supported Collaborative Learning*, 10(4):395–423, 2015.
- [8] R. De Lisi and S. L. Golbeck. Implications of piagetian theory for peer learning. 1999.
- [9] R. Decker. Management team formation for large scale simulations. *Developments in Business Simulation and Experiential Learning*, 22, 1995.
- [10] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. Liblinear: A library for large linear classification. *The Journal of Machine Learning Research*, 9:1871–1874, 2008.
- [11] O. Ferschke, I. Howley, G. Tomar, D. Yang, and C. Rose. Fostering discussion across communication media in massive open online courses. In *CSCL*, 2015.
- [12] K. Ghadiri, M. H. Qayoumi, E. Junn, P. Hsu, and S. Sujitparapitaya. The transformative potential of blended learning using mit edx 6.002 x online mooc content combined with student team-based learning in class. *environment*, 8:14, 2013.
- [13] G. Gweon. *Assessment and support of the idea co-construction process that influences collaboration*. PhD thesis, Carnegie Mellon University, 2012.
- [14] G. Gweon, M. Jain, J. McDonough, B. Raj, and C. P. Rosé. Measuring prevalence of other-oriented transactive contributions using an automated measure of speech style accommodation. *International Journal of Computer-Supported Collaborative Learning*, 8(2):245–265, 2013.
- [15] M. Joshi and C. P. Rosé. Using transactivity in conversation for summarization of educational dialogue. In *SLaTE*, pages 53–56, 2007.
- [16] R. F. Kizilcec, C. Piech, and E. Schneider. Deconstructing disengagement: analyzing learner subpopulations in massive open online courses. In *Proceedings of the Third International Conference on Learning Analytics and Knowledge*, pages 170–179. ACM, 2013.
- [17] R. F. Kizilcec and E. Schneider. Motivation as a lens to understand online learners: Toward data-driven design with the olei scale. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 22(2):6, 2015.
- [18] C. Kulkarni, J. Cambre, Y. Kotturi, M. S. Bernstein, and S. Klemmer. Talkabout: Making distance matter with small groups in massive classes. In *CSCW*, 2015.
- [19] I. Lykourantzou, A. Antoniou, and Y. Naudet. Matching or crashing? personality-based team formation in crowdsourcing environments. *arXiv preprint arXiv:1501.06313*, 2015.
- [20] S. MacNeil, C. Latulipe, B. Long, and A. Yadav. Exploring lightweight teams in a distributed learning environment. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education*, pages 193–198. ACM, 2016.
- [21] C. Piech, J. Huang, A. Nguyen, M. Phulsuksombati, M. Sahami, and L. Guibas. Learning program embeddings to propagate feedback on student code. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 1093–1102, 2015.
- [22] S. D. Teasley, F. Fischer, A. Weinberger, K. Stegmann, P. Dillenbourg, M. Kapur, and M. Chi. Cognitive convergence in collaborative learning. In *International conference for the learning sciences*, pages 360–367, 2008.
- [23] M. Wen, D. Yang, and C. P. Rosé. Virtual teams in massive open online courses. *Proceedings of Artificial Intelligence in Education*, 2015.
- [24] A. W. Woolley, C. F. Chabris, A. Pentland, N. Hashmi, and T. W. Malone. Evidence for a collective intelligence factor in the performance of human groups. *science*, 330(6004):686–688, 2010.
- [25] Z. Zheng, T. Vogelsang, and N. Pinkwart. The impact of small learning group composition on student engagement and success in a mooc. *Proceedings of Educational Data Mining*, 7, 2015.