

Informing the Design of Collaborative Activities in MOOCs using Actionable Predictions

Erkan Er
Universidad de Valladolid
Valladolid, Spain
erkan@gsic.uva.es

Eduardo Gómez-Sánchez
Universidad de Valladolid
Valladolid, Spain
edugom@tel.uva.es

Miguel L. Bote-Lorenzo
Universidad de Valladolid
Valladolid, Spain
migbot@tel.uva.es

Juan I. Asensio-Pérez
Universidad de Valladolid
Valladolid, Spain
juaase@tel.uva.es

Yannis Dimitriadis
Universidad de Valladolid
Valladolid, Spain
yannis@tel.uva.es

ABSTRACT

With the aim of supporting instructional designers in setting up collaborative learning activities in MOOCs, this paper derives prediction models for student participation in group discussions. The salient feature of these models is that they are built using only data prior to the learning activity, and can thus provide actionable predictions, as opposed to post-hoc approaches common in the MOOC literature. Some learning design scenarios that make use of this actionable information are illustrated.

ACM Classification Keywords

Collaborative learning; e-learning; supervised learning by classification

Author Keywords

MOOC; actionable learning analytics; collaborative learning.

INTRODUCTION

Massive Open Online Courses (MOOCs) offer quality online learning materials, reaching out to a wide audience. Nevertheless, the massiveness of the learner populations has reinforced the use of cognitivist-behaviorist pedagogies [9], which rely highly on information transmission and do not promote social learning. Although discussion forums and peer reviews are commonly incorporated to foster collaboration, being mostly implemented in an unstructured way with limited instructor facilitation, their effect on peer interaction and social learning has been limited [11]. Consequently, peer interaction and support barely exist in MOOCs [9].

Recently, collaborative learning within small teams sharing a common learning goal has been tested in MOOCs to promote

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

L@S 2019, June 24–25, 2019, Chicago, IL, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 123-4567-24-567/08/06...\$15.00

DOI: http://dx.doi.org/10.475/123_4

engagement and achievement of learners [15, 12]. In the research on supporting collaborative learning in MOOCs, most works have focused on offering technological solutions to perform group formation automatically [12, 15]. While they are valuable in demonstrating the effect of different criteria in different contexts, they do not support instructors to make informed decisions, and thus they continue to set group size and requirements based on their previous teaching experiences in formal non-massive online courses [12]. The prediction of student behavior has also been researched to estimate dropout, engagement and, sometimes, potential for collaboration, but most of the prediction literature uses post-hoc models that do not produce actionable information [6].

This paper presents the results of a work-in-progress that, using data gathered from two MOOCs, explored the use of transfer learning techniques, namely in-situ learning and transfer across courses [3]. These techniques allow generating predictions before the target label is produced (i.e., group discussion is completed) so that they are actionable for the upcoming activity. The predictions can thus be operationalized by instructors in several ways such as adjusting the group size, defining different roles (e.g., coordinator, spokesman), deciding on the workload on groups, etc.

BACKGROUND

Collaborative Learning in MOOCs

Though collaborative learning offers many benefits to students [4], its implementation in MOOCs is challenging. For example, manual group formation, carried out by the instructors, is not feasible. Moreover, the large variety within the MOOC learners' characteristics (e.g., learning preferences, educational background) and goals (e.g., aiming for the certificate, only intending to get familiar with a specific topic) can also hinder collaboration in groups [1], but instructor support and interaction with students is unfeasible at a massive scale [9]. Therefore, in most MOOCs, open discussions and peer reviews with neither a specific structure nor frequent instructor interference have been used optimistically to create collaboration [11].

To support instructional designers in introducing collaboration, the literature has explored algorithmic approaches to group formation that can outperform random pairing. It has been found that groups formed by clustering of students using certain criteria yield less dropouts [15] and more social interactions [12] compared to the groups formed by random matching. [13] set up groups using theory-driven team formation principles (e.g., heterogeneity in knowledge supports creativity in the group), achieving more productive engagement. Nevertheless, all these approaches for collaborative learning support are configured using instructional design decisions taken beforehand: instructors are not provided with early actionable information useful to make design decisions about the collaborative activity while the course is running

MOOC Research on Generating Actionable Predictions

We identify two important gaps in the MOOC prediction literature. First, the prediction research in MOOC contexts so far has mainly focused on dropout prediction [8], overlooking other essential learning behaviors such as engagement in collaborative learning activities. Machine learning predictions about the participation levels of students in collaborative activities may be actionable by instructors in several ways, such as determining an appropriate group size or determining the phases of the collaboration.

Second, existing models afford very limited use by instructors in real-world practice. According to a recent literature review [8], fewer than 10% of the surveyed proposals are feasible for real world as they are usually built on posterior data (which is available only after the target action is realized). Two alternative techniques to post-hoc approaches are transfer across courses (which involves the use of a model trained with past course data to make predictions in an ongoing course), and in-situ learning (in which student engagement in a past activity is used to generate proxy labels for training in the same course). So far, the research has noted promising results regarding the use of transfer learning for producing accurate predictions that are actionable for real-world use [5, 14, 2]. However, to the best of our knowledge, there are no previous studies on producing actionable information for supporting the design of collaborative learning in MOOCs.

THE STUDY

In this work, we applied transfer across courses and in-situ learning using data from two runs of the same MOOC to generate actionable predictions regarding students' participation in group discussions. The context of the study and the methods to build and assess the predictive models are described next.

Context

The context was a MOOC about translation of business terms between English and Spanish, offered in the Canvas Network MOOC platform, lasting seven weeks. Data from two runs of the course, in 2017 and 2018, were used in this research. In the first run, 1031 students enrolled, 668 of whom answered an initial questionnaire, and finally 161 completed the mandatory activity of the last week (hence, we consider the rest to have dropped out). The figures for the 2018 run were 866, 648 and 186, respectively.

	2017		2018	
	1 st activ.	2 nd activ.	1 st activ.	2 nd activ.
groups	160	150	108	118
active students	151	145	189	174
passive students	244	265	376	404

Table 1. Group activity information for both courses

Two collaborative activities (one in the fourth week and another in the sixth week) were designed, involving writing a report in teams that includes listing some terms extracted along with translations, providing the rationales for the terms, and citing the resources used. The instructional designers decided that they would be carried in groups of six. Table 1 shows the number of collaborative groups, and active (who did post in the discussion group) and passive students (who did not) in these groups. It can be seen that passive students are many more than active ones. Also there were less active students in the second activities.

Method

The prediction task in this study is to identify if students would post in their group discussions or not. Below, the training techniques applied, features generated, and the assessment metric of model performance are described.

Model Training Techniques

Three techniques were applied to build predictive models: post-hoc, in-situ learning, and transfer across courses. Logistic regression (in the Python implementation available in the *Scikit Learn* framework) was used as the classifier as it has shown to be effective in other tasks in the MOOC literature [14, 2].

To build post-hoc models, the dataset includes the labels (i.e. the student did or did not post) which are only available *after* the collaborative activity has finished. 10-fold cross validation is then used to train and test the model (i.e. the data is split in 10 folds; 9 are used for training and 1 for test; the process repeated for each combination of folds, and results averaged). This post-hoc model represents an optimistic reference point to compare the accuracies of the other methods [3, 14, 2].

In-situ learning is a transfer learning technique that involves creating a proxy label based on student engagement in a past learning activity and training a model using the data with proxy labels [3]. In this way, a trained model could be used to make a prediction about student engagement or success for an upcoming learning activity. In the current work, we apply in-situ learning to make predictions for both collaborative activities but using a different proxy label in each activity. For the first collaborative task, proxy labels were generated according to student participation in a different learning activity that was completed one week before. This activity was a mandatory assignment that involved analysis and revision of a given financial text. For the second collaborative task, student participation in group discussions of the first collaborative activity was used to create proxy labels. Note that this allows to train models *before* the collaborative learning starts, so that predictions can be operationalized by the instructor.

Feature	Description
x1	Total number of discussion posts
x2	Total number of quiz attempts
x3	Total quiz score
x4	Total time spent on quizzes
x5	Total number of assignment submissions
x6	Total page views
x7	Total views on introduction pages
x8	Total views on lecture pages
x9	Total views on discussion forums
x10	Total views on quiz pages
x11	Total views on assignment pages
x12	Total views on review-video pages

Table 2. Features used to train the predictive models.

Course	Training approach	1 st activity		2 nd activity	
		AUC	κ	AUC	κ
2017	post hoc	0.85	0.70	0.93	0.84
	in-situ learning	0.84	0.69	0.92	0.82
2018	post hoc	0.88	0.72	0.96	0.88
	in-situ learning	0.88	0.71	0.93	0.80
	transfer across course	0.87	0.70	0.94	0.82

Table 3. The AUC and Cohen’s κ scores of prediction accuracies.

Transfer across courses involves the use of past course data to train models that can be applied in a future course to make predictions [3]. In the current study, two models (one for each collaborative activity) were trained using the data from the course taught in 2017. Then, these models were applied in the 2018 version to predict student engagement in the first and the second collaborative activities respectively.

Feature Generation

To build the prediction models, 12 features were generated as listed in Table 2. These features account for student engagement in various course activities such as quizzes and discussions. These features are found to have predictive power in related tasks in the MOOC literature [2, 6].

Model Assessment

The performance was measured using area under the curve (AUC) and Cohen’s Kappa, which are metrics often used for imbalanced class distributions [10]. AUC scores above 0.8 or κ above 0.75 are considered from very good to excellent.

RESULTS AND USAGE SCENARIOS

In the first course, four models were built, using the post-hoc and in-situ approaches for each of the two collaborative learning activities. In the second run six models were derived, using post-hoc, in-situ and transfer across courses, for the two activities. Their performances are given in Table 3.

In both runs of the course, in-situ models yielded accurate results in predicting student participation in collaborative activities, with scores very close to the post-hoc models (an optimistic approach unfeasible in real-world [3]). Similar results were reported in the literature for predicting dropout

Feature	2017		2018	
	1 st activ.	2 nd activ.	1 st activ.	2 nd activ.
x1	0.00	0.13	0.00	0.00
x2	0.00	0.00	-0.05	0.00
x3	0.45	0.31	1.33	0.19
x4	-0.05	-0.19	0.00	-0.05
x5	0.76	1.06	-0.42	1.42
x6	0.00	0.00	0.00	0.00
x7	-0.07	0.00	-0.00	-0.35
x8	0.00	0.02	0.00	0.00
x9	0.20	0.02	0.00	0.00
x10	0.00	0.00	-0.39	0.00
x11	1.21	1.57	1.97	3.05
x12	0.00	0.00	0.36	0.00

Table 4. Coefficients of the features in the decision function (from Scikit LogisticRegression) of the four in-situ predictive models (transfer across courses uses the 2017 models to predict in 2018).

	1 st activity	2 nd activity
At least 2	46/108	53/118
At least 3	60/108	68/118
At least 4	64/108	77/118

Table 5. Number of at-risk groups in both collaborative activities based on three different thresholds.

[14] and video or exercise based engagement [2]. Moreover, the transfer across courses technique also reached accurate predictions for the second run of the course, in line with the results reported in [14]. These results show that transfer learning techniques can be used to train machine learning models without labels yet unavailable. All trained models performed better in the second collaborative activity, which can be due to the fact that MOOC participants with inconsistent behavior and activities (whose behavior is more difficult to predict) are likely to drop out of the course in later phases. Another potential explanation of this increase in the case of in-situ models comes from the fact that the behavior in the first collaborative activity makes a better proxy label for the second.

In order to better understand the derived models, their regression coefficients are presented in Table 4, showing that students with better quiz scores, more submission page views and more assignments submissions were more likely to post in collaborative activities, i.e. in general more engaged students will collaborate more.

In summary, these models achieve accuracies similar to post-hoc models, but with the benefit of being timely available. It is then worth debating how these actionable predictions could be used by instructors when setting up collaborative learning activities.

Identifying groups at risk of disengagement

The developed models can serve to identify beforehand which student groups are at risk of disengagement, so that instructors can act to prevent their possible failure in collaboration [12]. Students’ individual probabilities of participation can serve to

estimate the group behavior, and a threshold (i.e. at least n students should post in groups of N) can determine which groups are at risk. To illustrate this idea in practice, for the second run of the MOOC we used the in-situ models to compute the probabilities of at least $n = 2, 3, 4$ students posting in groups of $N = 6$, and considered those with probability lower than 0.5 to be at risk.

The results collected in Table 5 show a high number of at-risk groups. This information could be actioned in several ways: if there were very few groups in this condition, instructors could reassign students to improve chances in all groups; if, like here, there are many at-risk groups, instructors may need to act differently, as discussed next.

Changing the activity requirements and timeline

Instructors may want to relax the activity requirements in a way that the expected participation matches the demanded workload. For example, in the current MOOC context, groups were formed with six members to work together to prepare a comprehensive report that includes 20 terms with translations, resources used, and the rationale for the chosen terms. With this information, some of this requirements could be lowered (e.g. including less terms or making explanations voluntary). This information also affects decisions of the timeline. In this course, the instructors made the reflection that one week should be enough for the task with six member groups, but if many groups are in risk, the deadline could be postponed.

CONCLUSIONS AND FUTURE RESEARCH

There is a literature body on predicting student performance in MOOCs, but mostly using post-hoc models and focusing on individual behavior. This paper has presented prediction models built using techniques that produce actionable information affecting collaborative learning activities in MOOCs. The results are promising, with accuracies similar to post-hoc models and many potential opportunities of application. The study needs to be broadened in several ways: building models on data of other courses; rebalancing the output classes and analyzing the influence of this, as [7]; testing these models in real educational practice; and refining them according to the lessons learned to provide more useful indicators (e.g. adding more social features, or producing a ranked prediction –intensive, moderate, little or no collaborator–).

ACKNOWLEDGMENTS

This research is partially funded by the ERDF and the National Research Agency of the Spanish Ministry of Science, Innovation and Universities under grants TIN2017-85179-C3-2-R and TIN2014-53199-C3-2-R, by the ERDF and the Regional Council of Education of Castilla y Leon under grant VA257P18 and by European Union's Horizon 2020 under the Marie Skłodowska-Curie grant agreement 793317.

REFERENCES

1. C. Alario-Hoyos, M. Perez-Sanagustin, C. Delgado-Kloos, H.A. Parada, and M. Muñoz Organero. 2014. Delving into participants' profiles and use of social tools in MOOCs. *IEEE Transactions on Learning Technologies* 7, 3 (2014), 260–266.
2. M.L. Bote-Lorenzo and E. Gómez-Sánchez. 2018. An approach to build in situ models for the prediction of the decrease of academic engagement indicators in Massive Open Online Courses. *Journal of Universal Computer Science* 24, 8 (2018), 1052–1071.
3. S. Boyer and K. Veeramachaneni. 2015. Transfer learning for predictive models in Massive Open Online Courses. In *Proc. Conf. Artificial Intelligence in Education*. 54–63.
4. P. Dillenbourg. 1999. *Collaborative Learning: Cognitive and Computational Approaches*. Elsevier Science.
5. M. Ding, Y. Wang, E. Hemberg, and U. O'Reilly. 2019. Transfer Learning Using Representation Learning in Massive Open Online Courses. In *Proc. Int. Conf. on Learning Analytics & Knowledge*. 145–154.
6. E. Er, E. Gómez-Sánchez, M.L. Bote-Lorenzo, J.I. Asensio-Pérez, and Y. Dimitriadis. 2018. Supporting group formation in ongoing MOOCs using actionable predictive models. In *Proc. 5th Learning with MOOCs*.
7. E. Farrow, J. Moore, and D. Gašević. 2019. Analysing Discussion Forum Data: A Replication Study Avoiding Data Contamination. In *Proc. Int. Conf. on Learning Analytics & Knowledge*. 170–179.
8. J. Gardner and C. Brooks. 2018. Student success prediction in MOOCs. *User Modeling and User-Adapted Interaction* 28, 2 (2018), 127–203.
9. K.F. Hew and W.S. Cheung. 2014. Students' and instructors' use of massive open online courses (MOOCs): Motivations and challenges. *Educational Research Review* 12 (2014), 45–58.
10. L.A. Jeni, J.F. Cohn, and F. Torre. 2013. Facing imbalanced data recommendations for the use of performance metrics. In *Proc. 5th Conf. on Affective Computing and Intelligent Interaction*. 245–251.
11. K. Manathunga and D. Hernández-Leo. 2015. Has research on collaborative learning technologies addressed massiveness? A literature review. *Educational Technology & Society* 18, 4 (2015), 357–370.
12. L. Sanz-Martínez, E. Er, A. Martínez-Monés, Y. Dimitriadis, and M.L. Bote-Lorenzo. 2019. Creating collaborative groups in a MOOC: a homogeneous engagement grouping approach. *Behaviour & Information Technology* (2019).
13. H. Spoelstra, P.V. Rosmalen, T. Houtmans, and P. Sloep. 2015. Team formation instruments to enhance learner interactions in open learning environments. *Computers in Human Behavior* 45, 2015 (2015), 11–20.
14. J. Whitehill, K. Mohan, D. Seaton, Y. Rosen, and D. Tingley. 2017. MOOC dropout prediction: How to measure accuracy?. In *Proc. Learning@Scale*. 161–164.
15. Z. Zheng, T. Vogelsang, and N. Pinkwart. 2015. The impact of small learning group composition on student engagement and success in a MOOC. In *Proc. 8th Int. Conf. on Educational Data Mining*. 500–503.