

Comparing Peer Recommendation Strategies in a MOOC

François Bouchet, Hugues Labarthe, Kalina Yacef, Rémi Bachelet

► **To cite this version:**

François Bouchet, Hugues Labarthe, Kalina Yacef, Rémi Bachelet. Comparing Peer Recommendation Strategies in a MOOC. UMAP'17 4th international workshop on Educational Recommender Systems, Jul 2017, Bratislava, Slovakia. pp.129-134, 10.1145/3099023.3099036 . hal-01531503

HAL Id: hal-01531503

<https://hal.archives-ouvertes.fr/hal-01531503>

Submitted on 12 Feb 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Comparing Peer Recommendation Strategies in a MOOC

F. Bouchet
Sorbonne Universités
UPMC, Univ. Paris 06
CNRS, LIP6 UMR 7606, France
francois.bouchet@lip6.fr

H. Labarthe
Sorbonne Universités
UPMC, Univ. Paris 06
CNRS, LIP6 UMR 7606, France
hugues.labarthe@ac-creteil.fr

K. Yacef
School of Information Technologies
The University of Sydney
Australia
kalina.yacef@sydney.edu.au

R. Bachelet
Centrale Lille
University of Lille
France
remi.bachelet@ec-lille.fr

ABSTRACT

Lack of social relationship has been shown to be an important contribution factor for attrition in Massive Open Online Courses (MOOCs). Helping students to connect with other students is therefore a promising solution to alleviate this phenomenon. Following up on our previous research showing that embedding a peer recommender in a MOOC had a positive impact on students' engagement in the MOOC, we compare in this paper the impact of three different peer recommenders: one based on socio-demographic criteria, one based on current progress made in the MOOC, and the last one providing random recommendations. We report our results and analysis ($N = 2025$ students), suggesting that the socio-demographic-based recommender had a slightly better impact than the random one.

CCS CONCEPTS

• **Human-centered computing** → **Collaborative and social computing**; *Collaborative and social computing theory, concepts and paradigms; Social recommendation* • **Human-centered computing** → **Collaborative and social computing**; *Empirical studies in collaborative and social computing* • **Human-centered computing** → **Human computer interaction (HCI)**; *HCI design and evaluation methods; User studies* • **Applied computing** → **Education**; *E-learning*

KEYWORDS

peer recommendation ; MOOC ; recommendation strategies ; attrition ; clustering

1 INTRODUCTION

Attrition in Massive Open Online Courses (MOOCs) is one of the major issues, and as such many studies have been undertaken in

order to predict it and understand which are the main factors leading to it [7, 9, 16]. One of the main findings of these works was that the lack of social relationship – either because one does not interact with the community or because one's contacts are themselves dropping out of the MOOC – is a major predictor of attrition [14]. Therefore, finding ways to encourage interactions between students seems like a promising approach towards attrition reduction, but also to improve the learning experience for all students [6]. Long before MOOCs existed, researchers found that social interactions are a key factor for successful learning [1], and this led to recommendations of practices to reinforce the sense of community in distance education [5, 15]. However, in a platform involving thousands of other students, how do students choose who to contact when facing an issue or simply wanting to chat about their learning experience? This is where peer recommender systems (PRS) can come into play: this specific kind of recommender systems [17], rather than recommending a learning resource or another class to follow [10, 13], suggests instead fellow learners to interact with. We previously showed the positive impact of integrating such a system into a MOOC [11, 12], but the question of which recommending strategy is best in such a context remains an open one. Three main approaches can be considered: (1) recommending similar peers, assuming that people are more inclined to make contact with a person who has similar background (same gender, geo-location, education level and prior MOOC experience) as opposed to someone very different; (2) recommending peers who have reached the same level in the MOOC, assuming those would be more likely to help if the goal is to contact them for some concrete question about the content of the course; (3) a random recommendation. In this paper, we use a control study ($N=2025$) to compare these three strategies, where the random recommendation was used as a baseline to compare the two other strategies against, in order to investigate whether they

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

UMAP'17 Adjunct, July 9-12, 2017, Bratislava, Slovakia
© 2017 Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-5067-9/17/07 \$15.00
<http://dx.doi.org/10.1145/3099023.3099036>

have a positive impact on student engagement in the MOOC. This engagement is measured in terms of attendance (i.e. how active one was on the MOOC platform), completion (i.e. how far one went in the MOOC), performance (i.e. how much one scored) and participation (i.e. how much one interacted with others).

The remainder of this paper is organized as follows: in section 2 we present the MOOC in which our PRS was embedded, its overall functioning, the experimental design and the data collected. In section 3 we compare the 3 recommending strategies considered (random, based on sociodemographic criteria and progress-based), before concluding with an analysis of the limits of this study and the perspectives this work opens.

2 A PEER-RECOMMENDING SYSTEM IN A MOOC

2.1 Experimental context

2.1.1 The experimental platform

Our PRS was deployed during the 6th session of a biannual French MOOC on project management, hosted on a modified Canvas platform powered by Unow [2], which took place in Fall 2015 for 9 weeks. Among the particularities of this MOOC, students who register are invited to participate in a PreMOOC phase (week -3 to 0) in which they can perform self-assessment, introduce themselves in discussion threads and get familiar with the platform. The core phase of the MOOC then takes place (week 1 to 4 included), with a typical content made of lecture videos, assignments, self-evaluation quizzes, etc. Finally, during the remaining 5 weeks (week 5 to 9), the core MOOC remains open and students are expected to follow 2 specialisations modules among a dozen available and to take the final exams. Overall, $N_{tot}=24,980$ students registered to the MOOC, among which only 7,716 took the first quiz in week 1, and 4,757 took the final exam between weeks 5 and 9 – which means the attrition rate was of 81%, a figure typical of what is usually found in MOOCs [4].

2.1.2 The PRS and chat widgets

The recommendation system we designed (described in more details in [11]) associates two elements: the recommendation widget itself, which was always visible in the left navigational menu of the MOOC, and a chat system in the bottom right-hand corner of the screen, used to send instant messages to peers recommended by the recommendation widget (cf. Figure 1). The PRS widget itself displayed for lists of contacts: the suggested ones (green), the ones marked as favorite (orange) and the ignored ones (grey). Each list contained thumbnails showing other students through their name and photo (cf. Figure 1A). When mousing over a thumbnail (cf. Figure 1B), one could see the beginning of the students' biography (if any) and 4 icons allowing to send a private message to that student (through the MOOC messaging system), send an instant message through the chat widget, add the student as a favorite (to be able to contact them again later) or ignore the student (for her/him not to be offered as a suggestion anymore). The chat widget, minimized by default (cf. Figure 1C), had a blinking envelope icon and a sound played when a message was received. When clicking on it, the widget expanded giving access to

two tabs: in the first one, favorite contacts appeared and could be added to a chat to initiate a discussion with up to 6 of them. The second tab gave access to the list of previous chats one created or had been invited into; a click on one of these chats opened it and allowed to read messages in that chat thread as well as to post additional messages (cf. Figure 1D).

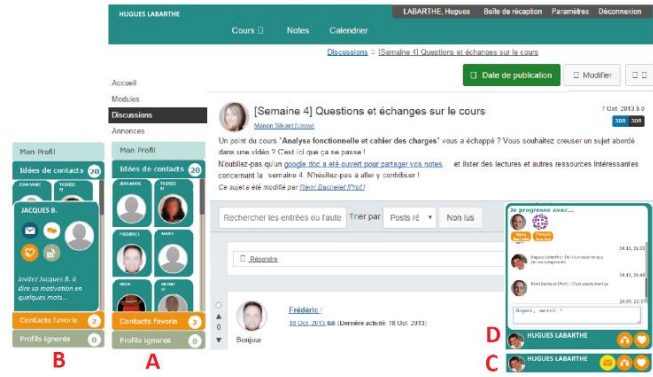


Figure 1: The PRS (left) and chat (right) widgets.

2.2 Experimental design and data collected

We conducted a control study, assigning students randomly to groups, for the purpose of evaluating the impact of a PRS on the student attrition rate [12]. The control group had no access to the PRS and the three PRS groups were, respectively:

- [Rnd] random recommendation,
- [Soc] recommendation based on a set of background-related criteria, including some general socio-demographic criteria (age, level of study, country, experience in MOOCs)
- [Prg] recommendation based first on progress in the MOOC (as indicated by the latest quiz one answered to) and then (for ties) on the socio-demographic criteria. Due to the dynamic nature of this type of recommendation, it was regularly recalculated (after each quiz) for each student.

To keep the four groups overall balanced, we assigned student to a group as they logged in on the MOOC for the first time after the deployment of the PRS: the first student was therefore in the control group, the second in [Rnd], etc. After demonstrating the positive impact of a PRS on student persistence and final grades in the MOOC [10], we focus in this paper on comparing the recommendation strategies. Therefore, we only consider the dataset containing the three treatment conditions ($N_{exp} = 6881$ students). Furthermore, for legal and ethical reasons, these students were first shown a short message asking them to accept the Terms of Use (ToU) of the experiment before they could see the PRS: $N_{ToU} = 2025$ students accepted these conditions and were therefore given access to the widgets. Among these, we considered in particular the students who also participated in the PreMOOC phase ($N_{ToU\&Pre} = 1417$), for reasons explained further on. In this subset, the number of students associated to each condition were respectively $N_{Rnd} = 296$, $N_{Soc} = 511$ and $N_{Prg} = 610$. The difference

between these three sample sizes may draw attention, however no experimental design factor can explain it: participation the PreMOOC happened before the deployment of the PRS, and the ToU shown to each of the three groups were identical. Across these three populations, $N_{int} = 272$ students interacted with both the PRS and the chat widgets (73 in [Rnd], 85 in [Soc], 114 in [Prg]).

3 COMPARING THE RECOMMENDERS

The research questions we address here are the following:

RQ1: “Were students in each recommender group similar before being given access to their PRS?”

Because of the experimental design we chose, where participants were randomly assigned to one of the 3 recommenders, our hypothesis was that there shouldn't be any difference between the 3 populations. This question underpins the soundness of our comparisons. Although students were randomly assigned to groups, we only retained those who accepted the ToU. In order to prove that there is no bias in our population, we need to show that these student populations are comparable. To do so, we performed inferential statistics tests on the sociological variables we had access to for our samples, and we verified that the populations were similar in terms of activity during the PreMOOC. This will be examined in section 3.1.

Then, through three research questions, we aimed to measure the differences of impact of the three recommenders on the students who had access to it.

RQ2: “Did more students interact with one of the PRS compared with the others?”

Our hypothesis was that a smarter recommendation algorithm could lead to an increase use of the associated recommender as students would be more encouraged to interact with people that are like them and/or have the same level of progression in the MOOC. We investigate this question by comparing the percentage of use of the 3 recommenders (section 3.3).

RQ3: “Among students who actually interacted with the PRS, did one of the recommenders have a stronger impact than the others in terms of attendance / completion / score / participation?”

Our hypothesis was that students who interacted with the *Soc* and *Prg* recommenders would interact more and do better than those who were provided with random recommendations. This will be examined in section 3.3

RQ4: “Among the students who attended the PreMOOC and then accepted the PRS terms of use, did one of the recommenders have a stronger impact than the others in terms of attendance / completion / score / participation?”

We already demonstrated in a previous paper [12] that the mere presence of the recommender had a positive impact on students' persistence and final grade in the MOOC. This question RQ4 addresses whether a particular recommender had a stronger impact. We do so by comparing the evolution of students between the

PreMOOC session and the MOOC session to see whether students who had access to a particular recommender did better than those who had access to the others. This will be examined in section 3.4.

3.1 Were students in each recommender group similar before being given access to their PRS?

In order to assess the similarity between the 3 groups of students who interacted with the three different recommender systems, we compared their social and behavioral features. In particular, we considered the 6 following variables: gender, country (grouped per sub-continent), previous experience in MOOCs, participation to the PreMOOC, age (discretized into 7 levels from “below 20 years old” to “above 70 years old”) and level of study (5 levels from “high school diploma” to “PhD level”). For the first 4 nominal variables, we performed 4 Chi-square tests which revealed no statistically significant differences between the 3 groups for gender ($\chi^2(2) = 1.45, p = .482$), experience in MOOCs ($\chi^2(2) = 1.31, p = .519$), participation to the PreMOOC ($\chi^2(2) = 2.08, p = .353$) and country ($\chi^2(24) = 26.07, p = .350$). For the last 2 ordinal variables, we performed 2 Kruskal-Wallis H tests which revealed no statistically significant differences for age ($\chi^2(2) = 0.34, p = .843$) and level of study ($\chi^2(2) = 1.44, p = .488$).

We can therefore conclude that our hypothesis was verified and that there were indeed no differences between the student populations exposed to the 3 different recommenders. Hence, it is safe to assume that any difference observed in the use of the recommender would be related to the recommender and not to intrinsic differences in students.

3.2 Did more students interact with one of the PRS compared with the others?

To answer RQ2, we compared the proportion of students in each group who interacted with their respective recommender. The results summarized in Table 1 reveal that there were no significant differences in the use of the recommender, regardless of the recommendation strategy used – which led us to reject our initial hypothesis. However, this observation means that any effect measured in further analyses would be exclusively related to the quality of the recommendation for students who used them, and not to the fact that more students used the recommendation in one group than in others.

Table 1: Percentage of activity with the recommender for each population

	Rnd	Soc	Prg
Did not accept the ToU	73%	69	71
		%	%
Accepted the ToU but did not interact with the recommender	23%	28	25
		%	%
Accepted the ToU and interacted with the recommender	4%	3%	4%

3.3 Among students who interacted with the PRS, did one of the recommenders have a stronger impact than the others?

In other words, among students who actually interacted with the PRS, was there a group doing better than the others in terms of attendance / completion / score / participation? To answer this, we compared (for the dataset made of $N_{ToU\&Pre}=1417$ students) each group with regards to various indicators of attendance, completion in the MOOC, scores and participation, which are detailed in Table 2. Raw results are shown in Table 3 and we used inferential statistics tests for each indicator. For indicators 1, 2, 3, 6, 8 and 10, we performed six one-way ANOVAs. Each distribution was preliminary tested for normality using the Shapiro-Wilk test: although they failed the test ($p < .05$ in each case), one-way ANOVA is considered robust against normality assumption and the homogeneity of variances criterion was always met ($p > .05$ using Levene's test). Outliers were preliminarily removed from each distribution according to their modified z-score value [8]. None of the results came out as statistically significant ($p > .10$ in all cases). For indicators 4, 5, 7 and 9, which are ordinal variables, we performed 4 Kruskal-Wallis H tests which revealed no statistically significant differences for the number of quiz attempts ($\chi^2(2) = 1.82, p = .403$) and the number of quiz completion ($\chi^2(2) = 0.71, p = .702$), but a statistically significant difference for the number of forum posts ($\chi^2(2) = 7.00, p = .030$) and a marginally significant difference for the number of private messages sent ($\chi^2(2) = 5.70, p = .058 < .10$) suggesting that students who were in the [Prg] and [Soc] groups exchanged more chat messages and posted more on forums than those in the [Rnd] (cf. columns 7 and 9). This result therefore supports partially our initial hypothesis: the students did not perform better, but at least they were more socially involved with others.

Table 2: Percentage of activity with the recommender for each population

Category	Indicators
Attendance	1. Number of days the student visited the platform
	2. Number of pages the student accessed
	3. Time spent on pages the student accessed [max = 600s to not count idle time online]
Completion	4. Number of attempts to complete a quiz
	5. Number of quizzes completed
Scores	6. Final score [31 compulsory quizzes + grade obtained in the final exam]
Participation	7. Number of posts on discussions (forums)
	8. Average length of discussion posts
	9. Number of messages sent via the chat widget (private messages)
	10. Average length of private messages

Table 3: Average and standard deviation (in italics) of persistence indicators for the 3 recommenders (max in bold) (* $p < .10$, ** $p < .05$)

Indicators	1	2	3	4	5	6	7**	8	9*	10
Rnd $N=296$	25.1	490	55	37.9	23.9	41.9	1.96	129	0.67	399
	<i>15.6</i>	<i>355</i>	<i>64</i>	<i>24.8</i>	<i>12.3</i>	<i>30.2</i>	<i>4.8</i>	<i>297</i>	<i>2.41</i>	<i>397</i>
Soc $N=511$	25.6	528	56.9	38.1	23.8	42.3	3.37	217	1.47	375
	<i>16</i>	<i>517</i>	<i>72.6</i>	<i>24.5</i>	<i>12.3</i>	<i>30.8</i>	<i>10.8</i>	<i>255</i>	<i>4.92</i>	<i>330</i>
Prg $N=610$	24.4	490	54.6	35.9	23.8	40.9	3.36	225	1.35	389
	<i>14.5</i>	<i>350</i>	<i>64.8</i>	<i>23.8</i>	<i>12.1</i>	<i>30.1</i>	<i>17</i>	<i>235</i>	<i>5.26</i>	<i>317</i>

3.4 Did one of the recommenders have a stronger impact on students?

To answer this question, we sought to analyse how students evolved as they were exposed to a particular recommender. To do this, we first extracted, using clustering, (i) typical student activity patterns in the PreMOOC period (i.e. before they were exposed to any recommender), and (ii) typical student attendance, participation and scoring patterns in the Core-MOOC 4-week period (i.e. after the recommender systems were made available). Then we examined how each of the 3 recommender groups evolved from PreMOOC clusters into Core-MOOC clusters, to see whether students who had access to one of the recommenders were doing better than those who had access to the others.

Table 4: Actions by students in each PreMOOC cluster

Features (in seconds)	PreMooc	PreMooc	PreMooc	PreMooc
	_D	_C	_B	_A
browsing_homepage	21	48	149	411
browsing_announce- ments	1	4	15	81
browsing_assignment	4	14	48	210
browsing_discuss._topics	2	8	26	190
browsing_grades	1	3	11	30
browsing_modules	7	43	140	428
browsing_pages	0	1	6	8
browsing_quizzes	0	1	2	2
downloading_assignment	0	0	0	2
viewing_assignment	1	11	49	208
viewing_calendar_events	0	0	0	7
viewing_discuss._topics	13	82	226	857
viewing_grades	0	0	1	1
viewing_modules	0	7	24	65
viewing_pages	25	163	550	1472
viewing_profiles	0	1	2	37
viewing_quizzes	33	768	1167	1965

PreMOOC clustering: We used as features the times spent on each of the 17 possible actions in the MOOC (cf. column 1 in Table 4). We used k-means algorithm and extracted 4 clusters, which centroids values are shown in Table 4 and called A, B, C and D on

the basis of their time spent (A the most active and D the least active). Students in cluster A spent over 1h40 on the website viewing lessons, quizzes and discussions (sum of the mean values). According to its means values, the second cluster (B) spent less than 40 minutes, essentially in the quizzes area; in the third cluster, C, the time is even shorter and those in the last one, D, stayed less than 2 min on the website in total.

The number of students who attended the PreMOOC is 4095. Among those, as mentioned before, $N_{ToU&Pre} = 1417$ accepted the ToU. Since there is no reason to believe that students associated to a given recommender but who have never actually seen any recommendation were affected by them, in the following analyses, we focus only on this subset of 1417 students. Moreover, because cluster PreMOOC_A is made of only 42 students, we chose to ignore it in the following analyses as they represent a sample size that is too small to be significant.

Core MOOC clustering: To measure the evolution of students in each group, we performed 3 more clusterings with the k-means algorithm using 3 sets of indicators: the attendance (indicators 1 to 3), the participation (indicators 7 to 10) and the final score (indicator 6), which provided us again with 4 clusters each. We named these clusters using the same convention as for the PreMOOC clusters, with A corresponding to the most active students, and D to the least active ones. Results for these 3 clusterings are provided in Tables 5, 6 and 7. These tables compare the distribution of the 3 original clusters into the 4 new ones, separately for each group. Therefore, to interpret each table, one should compare the 3 lines for the 3 groups associated to a given PreMOOC cluster. For instance, we see in Table 6 that for PreMooc_B, although only 6.8% of the students who had access to the random recommendation were among the most active ones (cluster Par_A), this number increases to 22% of those who had access to the [Soc] recommendation, and 17% of those who had access to the [Prg] recommendation. Overall, we tend to systematically observe a “shift to the right” when comparing the distribution of the non-random recommendation to the other ones, *i.e.* students who were using a recommender including socio-demographic parameters were more likely to have a higher attendance, a higher score and participate more to forums than students who were using a recommender providing random recommendations.

For each variable (Att, Sco and Par), for each cluster made in the PreMOOC (D, C and B), we performed a Chi-square test to see if there was significant difference between the distribution of students between the 3 groups (Rnd, Soc, Prg). We only found one marginally statistically significant difference for attendance with cluster PreMooc_B ($\chi^2(6) = 11.04, p = .087 < .100$).

We also considered the subset of 289 students among the 1417 who at some point edited their photo or their biography, which may indicate a desire to be contacted by other students. We replicated the same analyses done on the sample of 1417 students, and found a marginally statistically significant difference for score in

cluster PreMooc_D ($\chi^2(6) = 11.30, p = .080 < .100$), and for participation in cluster PreMooc_B ($\chi^2(6) = 10.83, p = .094 < .100$). Although the results are only marginally significant and they appear only for some of the PreMOOC clusters, they tend to indicate that, despite the moderate use students made of the recommenders, providing them with a set of students similar to them was more efficient than simply providing them with a set made of any students to contact.

Table 5: Attendance: evolution of the learners from the PreMOOC to the Core-MOOC periods

From \ To	Att_D	Att_C	Att_B	Att_A	Group
PreMooc_D (N=655)	34.9%	10.3%	40.4%	14.4%	Rnd (n=146)
	39.9%	10.9%	38.7%	10.5%	Soc (n=238)
	41.7%	6.6%	39.9%	11.8%	Prg (n=271)
PreMooc_C (N=479)	30.9%	21.6%	36.1%	11.3%	Rnd (n=97)
	40.9%	20.5%	31.6%	7.0%	Soc (n=171)
	37.0%	21.3%	32.2%	9.5%	Prg (n=211)
PreMooc_B (N=241)	47.7%	18.2%	22.7%	11.4%	Rnd (n=44)
	34.1%	37.4%	9.9%	18.7%	Soc (n=91)
	39.6%	30.2%	18.9%	11.3%	Prg (n=106)

Table 6: Participation: evolution of the learners from the PreMOOC to the Core-MOOC periods

From \ To	Par_D	Par_C	Par_B	Par_A	Group
PreMooc_D (N=655)	5.5%	58.9%	32.2%	3.4%	Rnd (n=146)
	2.1%	63.9%	28.6%	5.5%	Soc (n=238)
	3.3%	67.9%	24.0%	4.8%	Prg (n=271)
PreMooc_C (N=479)	6.2%	73.2%	14.4%	6.2%	Rnd (n=97)
	5.3%	65.5%	21.1%	8.2%	Soc (n=171)
	2.8%	72.5%	15.2%	9.5%	Prg (n=211)
PreMooc_B (N=241)	4.5%	54.5%	34.1%	6.8%	Rnd (n=44)
	7.7%	39.6%	30.8%	22.0%	Soc (n=91)
	6.6%	51.9%	24.5%	17.0%	Prg (n=106)

Table 7: Final score: evolution of the learners from the PreMOOC to the Core-MOOC periods

From \ To	Sco_D	Sco_C	Sco_B	Sco_A	Group
PreMooc_D (N=655)	26.0%	54.8%	7.5%	11.6%	Rnd (n=146)
	25.2%	53.8%	6.7%	14.3%	Soc (n=238)
	22.9%	56.8%	8.5%	11.8%	Prg (n=271)
PreMooc_C (N=479)	6.2%	73.2%	10.3%	10.3%	Rnd (n=97)
	11.7%	67.8%	9.9%	10.5%	Soc (n=171)
	12.3%	67.3%	6.2%	14.2%	Prg (n=211)
PreMooc_B (N=241)	4.5%	77.3%	9.1%	9.1%	Rnd (n=44)
	3.3%	82.4%	5.5%	8.8%	Soc (n=91)
	4.7%	72.6%	6.6%	16.0%	Prg (n=106)

4 CONCLUSIONS, LIMITS & PERSPECTIVES

After establishing that a peer recommender system generally improves student engagement in a MOOC [10], we investigated in this paper the differences of impact between specific recommendation strategies. We found that the recommendation strategy using background information seemed moderately more efficient than the random one (baseline), and that the recommender based on progress was not more efficient than the baseline.

However, this result needs to be put in the context of our experiment in which students did not know which recommending strategy was used to recommend them other students. In particular, the only concrete way for students to notice at first that recommended students were similar to them was if they had edited their photos (which may reveal, partially, the location and age) and biography. Otherwise, it was not until they had successfully initiated contact with another student that they could potentially notice this. Due to the overall low number of interactions through the chat, we can therefore consider that many students were not really given the possibility to notice which features the algorithm was using to recommend them other students. It is even more critical in the case of the PRS based on progression, as it required the interaction with other students to be about the content of the MOOC, and a qualitative analysis of messages has revealed that most interactions indicated more an intention to get to know other students than to ask them concrete questions based on the content. Another limit to consider is that our choice for the set of parameters to use for the socio-demographic recommendation might not be optimal: it was mainly linked to the type of data students had to provide when registering to the MOOC; it is very possible that another set of parameters could lead to better results.

Some of these limitations led us to consider modifying our PRS strategies by offering the students to choose themselves which type of recommendation they wanted [3], assuming that a background-based recommendation is probably more relevant to get to know other students, whereas a progress-based recommendation would be more relevant to answer to concrete questions.

ACKNOWLEDGMENTS

This work was funded by the French Educational Board and by the Human-Centred Technology Cluster of the University of Sydney. We thank Unow for partnering with us and deploying our RS on their MOOC.

REFERENCES

[1] Bandura, A. 1971. *Social Learning Theory*. General Learning Corporation.
 [2] Bouchet, F. and Bachelet, R. 2015. Do MOOC students come back for more? Recurring Students in the GdP MOOC. *Proc. of the European MOOCs Stakeholders Summit 2015* (Mons, Belgium, May 2015), 174–182
 [3] Bouchet, F., Labarthe, H., Bachelet, R. and Yacef, K. 2017. Who Wants to Chat on a MOOC? Lessons from a Peer Recommender System. *Digital Education: Out to the World and Back to the Campus* (Madrid, Spain, May 2017), 150–159.
 [4] Clow, D. 2013. MOOCs and the Funnel of Participation. *Proceedings of the Third International Conference on Learning Analytics and Knowledge* (New York, NY, USA, 2013), 185–189.

[5] Croft, N., Dalton, A. and Grant, M. 2010. Overcoming Isolation in Distance Learning: Building a Learning Community through Time and Space. *Journal for Education in the Built Environment*. 5, 1 (Jul. 2010), 27–64.
 [6] Ferschke, O., Yang, D., Tomar, G. and Rosé, C.P. 2015. Positive Impact of Collaborative Chat Participation in an edX MOOC. *Artificial Intelligence in Education: 17th International Conference, AIED 2015, Madrid, Spain, June 22–26, 2015. Proceedings*. C. Conati, N. Heffernan, A. Mitrovic, and F.M. Verdejo, eds. Springer International Publishing, 115–124.
 [7] Gütl, C., Rizzardini, R.H., Chang, V. and Morales, M. 2014. Attrition in MOOC: Lessons Learned from Drop-Out Students. *Learning Technology for Education in Cloud. MOOC and Big Data*. L. Uden, J. Sinclair, Y.-H. Tao, and D. Liberona, eds. Springer International Publishing, 37–48.
 [8] Iglewicz, B. and Hoaglin, D. 1993. Volume 16: How to Detect and Handle Outliers. *The ASQC Basic References in Quality Control: Statistical Techniques*. E.F. Mykytka, ed.
 [9] Kizilcec, R.F., Piech, C. and Schneider, E. 2013. Deconstructing Disengagement: Analyzing Learner Subpopulations in Massive Open Online Courses. *Proceedings of the Third International Conference on Learning Analytics and Knowledge* (New York, NY, USA, 2013), 170–179.
 [10] Klačnja-Milićević, A., Ivanović, M. and Nanopoulos, A. 2015. Recommender systems in e-learning environments: a survey of the state-of-the-art and possible extensions. *Artificial Intelligence Review*. 44, 4 (Dec. 2015), 571–604.
 [11] Labarthe, H., Bachelet, R., Bouchet, F. and Yacef, K. 2016. Increasing MOOC completion rates through social interactions: a recommendation system. *Proc. of the European Stakeholder Summit on experiences and best practices in and around MOOCs* (Graz, Austria, Feb. 2016), 471–480.
 [12] Labarthe, H., Bouchet, F., Bachelet, R. and Yacef, K. 2016. Does a Peer Recommender Foster Students' Engagement in MOOCs? *Proc. of the 9th International Conference on Educational Data Mining* (Raleigh, NC, USA, Jul. 2016), 418–423.
 [13] Manouselis, N., Drachsler, H., Vuorikari, R., Hummel, H. and Koper, R. 2011. Recommender Systems in Technology Enhanced Learning. *Recommender Systems Handbook*. F. Ricci, L. Rokach, B. Shapira, and P.B. Kantor, eds. Springer US, 387–415.
 [14] Rosé, C.P., Carlson, R., Yang, D., Wen, M., Resnick, L., Goldman, P. and Sherer, J. 2014. Social Factors That Contribute to Attrition in MOOCs. *Proceedings of the First ACM Conference on Learning @ Scale Conference* (New York, NY, USA, 2014), 197–198.
 [15] Rovai, A.P. 2002. Building Sense of Community at a Distance. *The International Review of Research in Open and Distributed Learning*. 3, 1 (Apr. 2002).
 [16] Yang, D., Sinha, T., Adamson, D. and Rosé, C.P. 2013. Turn on, tune in, drop out: Anticipating student dropouts in massive open online courses. *Proceedings of the NIPS Data-Driven Education Workshop* (2013).
 [17] Zheng, Y. and Yano, Y. 2007. A framework of context-awareness support for peer recommendation in the e-learning context. *British Journal of Educational Technology*. 38, 2 (Mar. 2007), 197–210.