Daniel Schiffner (Hrsg.): Proceedings of DeLFI Workshops 2018
co-located with 16th e-Learning Conference of the German Computer Society (DeLFI 2018)
Frankfurt, Germany, September 10, 2018

On the Emergence of Typical Behaviours in LMS

Truong-Sinh An[1], Christopher Krauss[2], Agathe Merceron[1]

Abstract: The emergence of Massive Open Online Courses (MOOCs) has enabled new research to analyze typical behaviours of learners. Inspired by this research, we characterize individual learning behaviours, taking into account specificities of the LMS we use. We then apply clustering techniques to uncover typical behaviours in university courses. In this contribution, we consider a classical face-to-face course on Advanced Web Technologies (AWT) delivered in a Master degree. This course has been offered in the winter semesters 2016/17 and 2017/18. Three typical behaviours appear in each course, reminiscent of those found by other researchers in MOOCs. Aggregating the data week by week, we investigate when these typical behaviours emerge. It turns out that they emerge only shortly before the exam for the two instances of the AWT course. We discuss implications of these findings.

Keywords: Learning Management Systems (LMS), Online Course, Typical Behaviors, Clustering.

1 Introduction

The emergence of MOOCs with the general observation of their low completion rates has triggered new research to analyze typical behaviours of learners in MOOCs. This brought forth evidence for various engagement/disengagement patterns as proposed by Kizilcec et al. [KPS13] or Ferguson & Clow [FC15]. Inspired by this research, we have investigated whether typical engagement patterns can be found in two courses backed by a learning management system (LMS) without being MOOCs, Java-FX, an optional online course in a Bachelor program and "Advanced Web Technologies" (AWT), a regular course in a Master program. As reported in [AMK17], we have found the following typical learning behaviors: (i) *completing*: students who have completed correctly most of the exercises offered in the course, (ii) *auditing*: students who did exercises infrequently, if at all, but consulted some other material, (iii) *disengaging*: students who solve exercises in the first learning unit of the course and they are not active anymore, and (iv) *weak completers*: students who do a number of exercises but not as many as those of the completing group. The *completing* group has been found in the two above mentioned courses and also in [KPS13, FC15]. The *auditing* group found also in the two courses bears similarities with the *auditing* and *sampling* group found in [KPS13] and with the *samplers* group found in [FC15]. The *disengaging* group found in the Java-FX course only reminds of the

[1] Beuth Hochschule für Technik Berlin, Luxemburger Straße 10, 13353 Berlin, {vorname}.{nachname}@beuth-hochschule.de
[2] Fraunhofer FOKUS, Kaiserin-Augusta-Allee 31, 10589 Berlin, {vorname}.{nachname}@fokus.fraunhofer.de

disengaging group of [KPS13] and *strong starters* of [FC15], while the *weak completing* group is specific to the AWT course.

In order to plan some intervention, it is essential to know when these behaviours emerge. In this contribution, we continue the work presented in An et al. [AMK17] by considering one more instance of the AWT Master course and by investigating how these typical behaviours emerge over time in this course.

This paper is organized as follows. Related works are discussed in Section 2. The AWT course and logged data are introduced in Section 3. Subsequently, typical learning behaviours are presented, their emergence is analyzed and results are discussed. Conclusion and future works are given in Section 5.

2 Related works

Some researches document very high drop-out rates in MOOCs [K14], up to 87% [OSB14]. Thereby, the analysis of related work indicates that the course success rate decreases with higher participation numbers. The reasons are, among others, low intention to complete the course, missing time, course difficulty and a lack of support [OSB14]. While dropout users represent a significant class of users due to their percentage and their negative impact, they represent just one type of learners. The rest of participants are users who engage with the course materials until they reach the course objective.

Kizilcec et al. [KPS13] investigated learners' engagement in MOOCs, which offer weekly videos and assessments and proposed four typical engagement/disengagement patterns as described in the introduction. These categories have been identified in three courses. Their proportions differ in each course. To discover these categories, they have first characterized a student by a tuple giving the status each week: "*on track* [T] (did the assessment on time), *behind* [B] (turned in the assessment late), *auditing* [A] (didn't do the assessment but engaged by watching a video or doing a quiz), or *out* [O] (didn't participate in the course in that week)" [KPS13].

In an attempt to replicate the work of Kizilcec et al. [KPS13], Ferguson and Clow [FC15] suggest that the methodology used to uncover typical learning behaviours in a course context does not necessarily generalize to another course adopting different elements of pedagogy and learning design. Since the MOOCs as analyzed by Ferguson and Clow [GRD16] follow a social constructivist pedagogy, they adopt the methodology of Kizilcec et al. [KPS13] by adding participation in forums.

Gelman et al. [GRD16] adopt a different, more bottom-up approach to discover typical behaviours in MOOCs: they use a set of 21 features that they can extract week by week from the log data and adapt non-negative matrix factorization to obtain weekly behaviours that are supported by a combination of those features. This approach is attractive because it does not need a manual selection of features to characterize the behaviour of a student; instead, the algorithm selects and combines features from the set it receives as input. A difficulty lies in the interpretation and the practical use of the discovered behaviours.

While an *auditing* behaviour is easy to comprehend [KPS13], it is less clear what a weekly *deep* behaviour "the associated students must have spent a long time on a single resource" as found in [GRD16] means for educators.

Graf and Kinshuk [GK08] study the behaviours of students in LMSs differently. Their aim is not to find typical engagement behaviours in the usage data. They assume that the learning style of a student is known. In a study involving 43 students, they have found that students with different learning styles do navigate differently through the resources of the course.

These works use the accumulated data at the end of the course and do not investigate when the behaviours emerge. McBroom et al. [MJK16] investigate the behaviours of computer science students in an auto-grading system. They have found six different types of submissions like early, normal, late and so on. They also studied the behaviours of the students regarding the types of submissions they do. They have found that, since the middle of the course, students tend to adopt the same submission type.

3 Structure of the courses

The courses have been created and taught with the smart learning infrastructure from the project "Smart Learning – digital media in vocational training"[3] funded by the German Ministry of Education and Research. The infrastructure is comprised of a learning management system, the learning companion app, a repository for learning objects, a recommendation engine and a learning analytics module [KMA17]. A course is essentially a sequence of learning units and a learning unit contains primarily learning objects. Each learning object is paired with its metadata that includes at least one learning objective. A learning object (LO) can be a piece of text including programming examples, a video, an exercise (similar to an exercise of an assessment in a MOOC), an animation and so on. The learning objectives of a learning unit are the union of the learning objectives of its learning objects. When opening a learning unit, the top item that can be opened is the list of the learning objectives of that unit. Learners can rate how much they know each learning objective, from 1 "hardly know anything" to 5 "expert". We call this list *self-assessments*. The next item after the sequence of LOs is again the list of learning objectives. By rating them, students can reflect on how much they know after learning the unit. Finally, a feedback item and a forum item complete any learning unit. Apart from its sequence of learning units, a course contains a schedule which specifies dates for the start and end of the course, as well as when each learning unit should be learned. All users' interactions are stored using the xAPI specification in the open-source learning record store called Learning Locker. All the learning material of the course AWT was available from the start of the course to encourage self-pacing and self-organization of students. Furthermore, the time schedule is not compulsory. There is no penalty if someone does not follow the schedule.

[3] https://projekt.beuth-hochschule.de/smart-learning/

The course Advanced Web Technologies (AWT) targets master computer science students. Technical experts taught in 12 presence lectures diverse topics that are of interest for future web developers – from web technology basics, such as HTML, over media delivery and content protection, to personalization through recommender systems and the Internet of Things. The lectures are mostly held with PowerPoint slides showing definitions, specifications, and source code, animations for concepts and videos for practical examples. The about 1000 presented slides are converted to digital learning objects, one slide being a single LO, and grouped into 105 learning units – with videos, animations and additional multiple-choice questions at the end of the learning units. Moreover, as some students still want to learn with a printed version of the slides, the last LO of a learning unit consists of a PDF file that can be downloaded and which contains all the slides of the unit; accordingly, the metadata of this LO corresponds to the totality of all the other LOs within this learning unit. At the end of the course, students can earn credits by completing a one-hour presence exam consisting of 50 multiple choice questions and five bonus questions.

A total of 142 students initially enrolled for AWT in winter semester 2016/17. However, there were 43 no-shows; "people register but never log in to the course while it is active" [H13]. Only the remaining 99 students are considered for the analysis in this paper. 75 students completed the final exam and the average grade was 1.90; only one student fail the exam. The users generated 92,825 xAPI statements in total during the 16 weeks of the course.

In the winter semester 2017/18, 75 active users generated under the same conditions 121,445 xAPI statements during the 18 weeks of the course. 53 students completed the final exam with an average grade of 2.0 and again one student failed.

4 Methodology and results

In the following, we motivate the selected feature set and explain the clustering techniques we use. We show the typical engaging behaviours that we have found, investigate when they emerge and discuss the results.

4.1 Methodology

The courses do not have assessments with deadlines; there is a suggested timetable for the content and there is no penalty if students do not follow it. Further, all the learning material is available from the start of the course and the infrastructure tracks details of the behaviour of students at the level of the learning objects. In An et al. [AMK17] several feature selection methods have been investigated. For the AWT course, the method *assessment scores* was most appropriate. Therefore, this feature set is used in this contribution.

Assessment scores or performance on all assessments: A student is represented by a vector that has the size of all assessments; values are ratings given in all self-assessments and marks earned in all exercises; all values are scaled scores between 0 and 1. Two students are similar if they achieved similar scores on all assessments. For the course AWT, a student is represented by a vector 246 features made of 196 self-assessments corresponding to the same amount of learning objectives and 50 exercises. Values for the features vary because few students self-assess themselves. The average value for self-assessment features is around 0.02 while it is around 0.3 for exercise features (missing values are per student and feature are set to zero).

In a first step, we use all the data stored during the whole course. We found three clusters in each course. In order to plan some intervention, it is essential to know when these behaviours crystallize. To do so, we clustered the accumulated data week by week and inspect how clusters evolve. We used RapidMiner and applied the X-means clustering algorithm with Euclidean distance.

4.2 Results

Taking the full data of the first instance of the course, X-means returns 3 clusters, as illustrated in diagram 1 last column on the right: week 16. The column shows the size of the three clusters while the dots correspond to the average value of all exercises of the cluster centre. Consider the top cluster in green: it contains nearly 30% of the students, actually 28 students; the average score of the cluster centre on all exercises - the green dot of this column - is 0.83. Looking at the data, one notices that these students have engaged with some self-assessments and nearly all exercises. If one sorts the students according to the number of distinct exercises they have solved in the course, 25 of these students are the top 25. They have worked nearly all the exercises out, on average 42 out of 50, and solved almost all of them right, therefore we call this cluster *completing*. The final exam mark in this *completing* cluster reaches 1.50 on average, a better mark than the overall average of 1.90. The second cluster in red in the middle of the column consists of 10 students who provided a few self-assessments and answered about 30% - 50% of the exercises. Students in this cluster rated self-assessments in the first three units and worked out exercises but with not so good scores; the average score of the cluster centre on all exercises - the red dot of this column - is 0.39. To some extent, they exhibit some kind of *completing* pattern in terms of exercises, because they solved almost half of them: on average 22 from a total of 50. Their average mark in the final exam is 2.03 that is slightly less good than the general average of 1.90. We call this cluster *weak completing*. The remaining students in the last cluster (bottom blue part of the column) have engaged in self-assessments and exercises sporadically all over the course and they did exercises infrequently if at all: on average 1 out of 50. However, they did access .pdf files. We call this cluster *auditing*. All learners who did not participate in the final exam fall into this cluster. The average score of the cluster centre on all exercises - the blue dot - is 0.02 and the average mark of the students in this cluster who participated in the final exam is 2.23, which is below the general average.

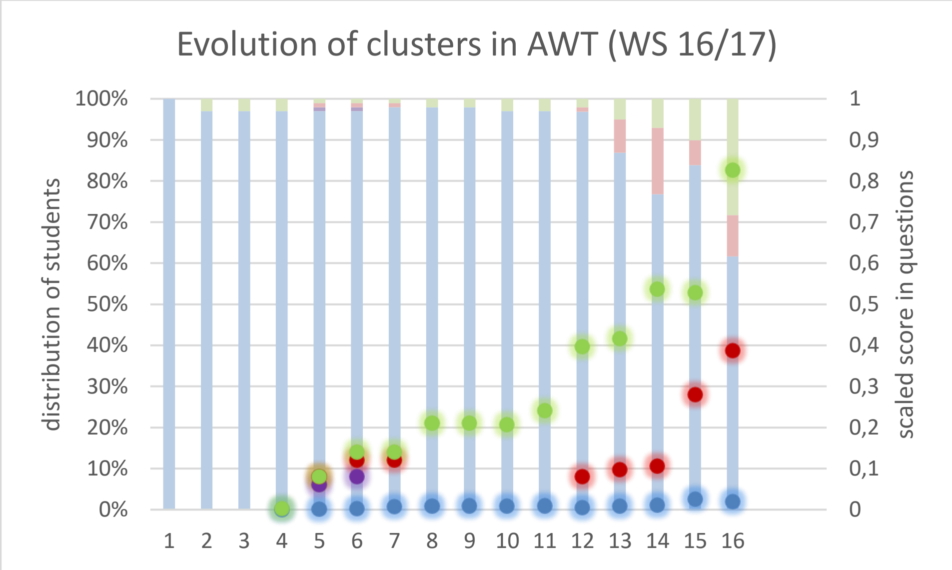

Diagram 1: The evolution of the clusters over the course AWT in winter semester 2016/17.

Let us look now at how the clusters form, from week 1 to week 15. In week 1, column 1 in Diagram 1, there is only one cluster. In week 2 and 3, the small cluster in green on the top of the column consists of three students who begin self-assessing themselves but do not solve exercises. That is why no dots in blue or green appear in the columns. Note that week 3 corresponds to a deadline in the time schedule. In week 4, these three students start solving exercises; see the green dot at the bottom of column 4 that gives the average score of the cluster centre. Two of these three students belong to the final *completing* group and one to the final *auditing* group. From the deadline in week 5 till the deadline in week 11, X-means isolates 3 than 2 then 1 small clusters of students who go ahead solving exercises; these students mostly belong to the final *completing* group. From the 12th week all three clusters appear; while students in the green cluster will mainly belong to the final *completing* cluster, students in the red cluster will mainly either belong to the *completing* or *weak completing* cluster. Students from the blue cluster changed their learning intensity and moved into the red or green cluster.

The results for the second instance of the course are similar to the results of the first instance, see Diagram 2. With all data, X-means returns three clusters see column 18. The three clusters have the same interpretation: *completing*, *weak completing* and *auditing*. They also have the same implication for the final exam: In the final exam, *completing* students obtained 1.30 in average, a better mark than the overall average of 2.00.

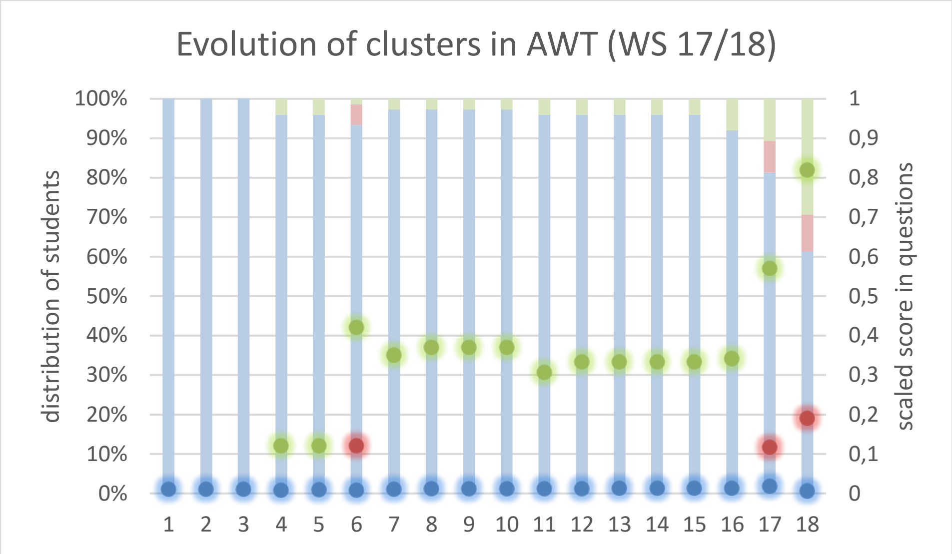

Diagram 2: The evolution of the clusters over the course AWT in winter semester 2017/18.

Students from the *weak completing* cluster obtained an average of 2.30, a lower mark than the overall average. Students from the *auditing* cluster who attended the final exam obtained an average of 2.70.

While clustering the data by week by week, one notices also similar trends compared with the first instance. Among the differences, clusters start to appear from week 4 but exercises are solved from the beginning; however, more engagement with exercises starts to really appear in week 17, green dot over 0.5, while it was in week 14 in the first instance.

4.3 Discussion

The clusterings obtained for the two instances of the AWT course show many similarities. This is in agreement with results obtained by Kidzinski et al. [KPS13]. Although their set-up is different from ours, the model Kidzinski et al. have obtained in one course generalizes well to another instance of the same course, but not to another course. Likewise, Ferguson and Clow [FC15] found similar engagement patterns in repeated courses.

A *completing* cluster is found in both courses. Such a cluster is also found in each course by Kizilcec et al. [KPS13] and Ferguson and Clow [FC15]. However, this cluster emerges only shortly before the end of the course in the two instances of the AWT course. One explanation might be the overall assessment of the AWT course. There is no mid-term assessment that counts for the final mark. On the contrary, the final mark is the one obtained in the final exam, which is composed of exercises similar to the ones appearing in the course material. Students might think that just-in-time learning is the optimal way to pass the exam. Students may earn an additional mark for realizing a student project covering the AWT topics. These projects are part of a subsequent course and can be assessed independently of the AWT course.

5 Conclusion and future work

We used X-Means clustering to extract typical behaviours of engagement in two instances of a face-to-face course in a master degree. In continuation of a previous work, we could put in evidence behaviours that remind of patterns found by Kizilcec et al. [KPS13]: *completing* and *auditing* and a third one *weak completers* in both analyses. In both instances, the students are increasingly concerned with exercises in the last weeks before the final exam. The *completing* students have solved nearly all exercises with a good mark, the *weak completing* students solved fewer exercises with not so good marks, and the *auditing* students did exercises infrequently if at all. This performance on exercises correlates with the mark of the final exam.

We presented the identified clusters to the instructors of AWT after the course. The most important finding is that the success in the final exam correlates with the number of activities on assessments and exercises. It is not clear, whether more learning activities in our system always lead to a better understanding and consequently to better marks, or students who wrote good exams would have been more active anyway (e.g., because they show a higher motivation or more interest in the topic). However, as a result of this research, the instructors want to offer a broader variety of exercises to the students in future instances of this course. Additionally, they wished a live visualization of the current clusters and their typical behaviours as a part of the teachers learning analytics dashboard.

Future work will use these results to incorporate gamification elements in those courses. To help improve students' performance, gamification elements should encourage students to solve the exercises of all learning units and also encourage them to solve the exercises correctly. Care has to be taken in defining those elements, as students could get the exercises right simply by attempting them till they find the right answer, thus gaming the gamification in some sense, and not learning anything.

The learning theory of goal-setting supports the self-regulation of students and ensures that students are aware of what is expected of them with their goal in mind. The theory of goal-setting assumes behaviour is a result of conscious goals and intentions, so students work towards their own objectives, gain self-satisfaction and motivation. Gamification elements help to implement this learning theory [LL91, G15]. Students could choose individually their own goals from a predefined set, e.g. "I want to solve more than half of all exercises by the end of the course". Gamification elements could be introduced with achievements by badges; once defined, badges would be automatically generated by the smart-learning infrastructure for each course. As an example, a badge could cover all exercises in one learning unit. The rules to win badges should encourage students to adopt a completing behaviour and do all exercises correctly similar to the learning theory of goal setting, e.g. a student receives the badge in gold if all exercises were completed correctly after at most two attempts, in silver after a maximum of 4 attempts and for more attempts in bronze. This encourages understanding of the material before an exercise is completed. The playful approach of achieving gold for all badges encourages the participants to solve all exercises carefully, not by chance. To encourage regular learning, badges could be linked to the time schedule of the course. Predefined goals are better under control of the entire smart learning infrastructure, which is transparently communicated with the students. With the option that students can change their goal during the course, unattainable goals would no longer be selectable. The transparent communication of possible goals shall encourage students to achieve goals with perseverance.

Acknowledgements

The authors would like to thank the whole Smart Learning team for their great work and many constructive ideas. We would also like to thank the instructors for the AWT course: Stephan Steglich, Louay Bassbouss, Stefan Pham, and André Paul. This work is partially supported by the German Federal Ministry of Education and Research grant number 01PD14002B and 01PD17002B.

References

[AMK17] An, T.-S., Merceron, A., Krauss, C., 2017. Can Typical Behaviours Identified in MOOCs be Discovered on other Courses? In In Proceedings of the 10th International Conference on Educational Data Mining, EDM’2017, Wuhan, China, June 25-28, p. 220-225.

[FC15] Ferguson, R. and Clow, D. 2015. Consistent commitment: Patterns of engagement across time in Massive Open Online Courses (MOOCs). Journal of Learning Analytics, 2(3), 55–80. http://dx.doi.org/10.18608/jla.2015.23.5.

[GRD16] Gelman, B., Revelle, M., Domeniconi, C., Johri, A., and Veeramachaneni, K. 2016. Acting the Same Differently: A Cross-Course Comparison of User Behavior in MOOCs. In Proceedings of the 9th International Conference on Educational Data Mining, EDM’2016, Raleigh, USA, June 29-July 2, p. 376-381.

[G15] Grund, C.K.: How Games and Game Elements Facilitate Learning and Motivation: A Literature Review. In: 45. Jahrestagung der Gesellschaft für Informatik, Cottbus, S. 1279-1293, 2015.

[GK08] Graf, S., Kinshuk. 2008. Analysing the behaviours of Students in Learning Management Systems with Respect to Learning Styles. In Advances in Semantic Media Adaptation and Personalization, (SCI) 93, Springer Verlag, p. 53–73.

[H13] Hill, P., Emerging student patterns in MOOCs: A (revised) graphical view, 2013. http://mfeldstein.com/emerging-student-patterns-in-moocs-a-revised-graphical-view/ accessed Sept. 26, 2017.

[KPS3] Kizilcec, R.F., Piech, C., and Schneider, E. 2013. Deconstructing disengagement: analyzing learner subpopulations in massive open online courses. In Proceedings of the Third international conference on learning analytics & knowledge LAK'13, New York, NY, ACM,170-179. doi>10.1145/2460296.2460330.

[K14] Kloft, Marius, et al. "Predicting MOOC dropout over weeks using machine learning methods." Proceedings of the EMNLP 2014 Workshop on Analysis of Large-Scale Social Interaction in MOOCs. 2014.

[KMA17] Krauss, C., Merceron A., An, T-S., Zwickelbauer M., Steglich, S., Arbanowski, S.: Teaching Advanced Web Technologies with a Mobile Learning Companion Application. Proceedings of the 16th ACM World Conference on Mobile and Contextual Learning (mlearn 2017), ACM, Ont. 30 - Nov. 1, 2017, Larnaca, Cyprus, doi>10.1145/3136907.3136937.

[MJK16] McBroom, J., Jeffries, B., Koprinska, I., and Yacef, K. 2016. Mining behaviours of students in autograding submission system logs. In Proceedings of the 9th International Conference on Educational Data Mining, EDM'2016, Raleigh, USA, June 29-July 2, p. 159-166.

[LL91] Latham, Gary P., Locke, Edwin A., Self-regulation through goal setting, Organizational Behavior and Human Decision Processes, Volume 50, Issue 2, 1991, Pages 212-247, ISSN 0749-5978, https://doi.org/10.1016/0749-5978(91)90021-K.

[OSB14] Onah, Daniel FO, Jane Sinclair, and Russell Boyatt. "Dropout rates of massive open online courses: behavioural patterns." EDULEARN14 proceedings (2014): 5825-5834.