



Open Research Online

The Open University's repository of research publications
and other research outputs

Examining engagement: analysing learner subpopulations in massive open online courses (MOOCs)

Conference Item

How to cite:

Ferguson, Rebecca and Clow, Doug (2015). Examining engagement: analysing learner subpopulations in massive open online courses (MOOCs). In: 5th International Learning Analytics and Knowledge Conference (LAK15), 16–20 March 2015, Poughkeepsie, NY, USA, ACM.

For guidance on citations see [FAQs](#).

© 2015 The Authors

Version: Version of Record

Link(s) to article on publisher's website:

<http://dx.doi.org/doi:10.1145/2723576.2723606>

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's [data policy](#) on reuse of materials please consult the policies page.

oro.open.ac.uk

Examining Engagement: Analysing Learner Subpopulations in Massive Open Online Courses (MOOCs)

Rebecca Ferguson and Doug Clow

Institute of Educational Technology

The Open University

Milton Keynes

MK7 6AA, UK

+44-1908-654956

firstname.surname@open.ac.uk

ABSTRACT

Massive open online courses (MOOCs) are now being used across the world to provide millions of learners with access to education. Many learners complete these courses successfully, or to their own satisfaction, but the high numbers who do not finish remain a subject of concern for platform providers and educators. In 2013, a team from Stanford University analysed engagement patterns on three MOOCs run on the Coursera platform. They found four distinct patterns of engagement that emerged from MOOCs based on videos and assessments. However, not all platforms take this approach to learning design. Courses on the FutureLearn platform are underpinned by a social-constructivist pedagogy, which includes discussion as an important element. In this paper, we analyse engagement patterns on four FutureLearn MOOCs and find that only two clusters identified previously apply in this case. Instead, we see seven distinct patterns of engagement: Samplers, Strong Starters, Returners, Mid-way Dropouts, Nearly There, Late Completers and Keen Completers. This suggests that patterns of engagement in these massive learning environments are influenced by decisions about pedagogy. We also make some observations about approaches to clustering in this context.

Categories and Subject Descriptors

K.3.1 [Computers and Education]: Computer Uses in Education; Distance Learning – *Massive Open Online Course, MOOC, Learner Engagement Pattern*

General Terms

Algorithms, Measurement

Keywords

Learning Analytics, Learner Engagement Patterns, MOOCs

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 ACM 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.

LAK '15, March 16 - 20, 2015, Poughkeepsie, NY, USA. Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3417-4/15/03...\$15.00

<http://dx.doi.org/10.1145/2723576.2723606>

1. INTRODUCTION

The generic name for MOOCs emphasizes commonalities of scale (massive), economic / philosophical perspective (open), location (online) and structure (course). At the same time, it omits a key area of difference, their underlying pedagogy, which can also be described as their approach to teaching and learning.

The original MOOCs, developed by Siemens and Downes, employed a connectivist approach to learning [5]. This approach is signaled by references to them as cMOOCs. Connectivism considers learning to be social, technologically enhanced, distributed within a network and associated with the recognition and interpretation of patterns. Knowledge is developed as a result of experience and the self-organizing nature of appropriately designed networks [5].

The cMOOCs were followed by xMOOCs. In this case, the defining letter 'x' did not refer to a specific pedagogy, but to the role of these MOOCs as an extension to a previous offering. However, in many cases, this 'extension' version took the essential elements of education to be content and assessment, with input from educators bundled as part of the content. This led to an instructivist approach to teaching and learning in which 'learning goals are predefined by an instructor, learning pathways structured by environment and learners have limited interactions with other learners' [12]. In Siemens' view, 'cMOOCs focus on knowledge creation and generation whereas xMOOCs focus on knowledge duplication' [15].

The FutureLearn platform takes a different approach. It employs a social-constructivist pedagogy, based on the Conversational Framework [11; 13]. This is a general theory of effective learning through conversations, with oneself and others, about the immediate world and about abstract concepts [14]. To engage in successful conversations, all parties need access to a shared representation of the subject matter as well as tools for commenting, responding and reflecting, and so these tools and shared representations formed part of the design of the FutureLearn platform.

Although these types of MOOC have different pedagogies, a problem that almost all of them experience is the large difference between numbers registering and numbers completing. Some of this drop-off can be explained in positive ways. MOOC registration can be regarded as similar to bookmarking a website

or placing an item temporarily in an online shopping basket – it does not necessarily represent a commitment to engage further. FutureLearn reports that around a third of those who register do not return to start the course, suggesting a more useful figure compares those who start and those who complete the course.

Another large group, characterized as ‘Samplers’, appears to be made up of those who take a brief look at a course, decide it is not for them, and then leave [9]. Respondents to surveys associated with FutureLearn MOOCs have shared various reasons for participating, including: to try out learning online, to find out more about MOOCs, or to find out more about a specific university. The learning objectives of these learners may be met without working through an entire course.

Another perspective is offered by Downes, who points out that most media are not consumed in their entirety by any single individual: ‘nobody thinks a library a failure if you don’t read everything in the collection, or an author a failure if you don’t read their entire corpus. And just so with MOOCs’ [6].

Despite these valid reasons and explanations for non-completion, the high drop-out rates on most MOOCs still provide a cause for concern. MOOCs are open, but ‘students seek not merely access, but access to success,’[4] and it seems that many encounter not success but failure when they study these courses. The steep drop-off of the ‘Funnel of Participation’ [3] is not welcomed by all.

Jordan’s work [7] provides an overview of this issue, and her website [8] provides a dynamic view of the data about MOOC completion rates that has been made public. The site currently presents data relating to around 180 MOOCs that are based on a variety of different platforms. The site deals with the percentage of registered learners who complete a course (rather than the percentage of people who actually started the course who go on to complete). It shows that MOOCs typically have completion rates of around 13%. At the time of writing, no MOOC with more than 60,000 registered students had reported a completion rate of higher than 13%, and only two MOOCs had achieved a completion rate of over 40%.

Some of the variability here is highly likely to be contextual. For example, it is easier to retain students on a two-week course than on a 12-week course. Nevertheless, the variability suggests that course context, course design and course pedagogy could have an effect on retention, and thus on students’ chances of success.

1.1 Using analytics to investigate MOOCs

Learning analytics are concerned with the use of trace data relating to learners and their contexts, for purposes of understanding and optimising learning and the environments in which it occurs [16]. They therefore offer a way of identifying factors that influence retention, thus enabling educators and platform providers to make changes to context, design and pedagogy where appropriate. The large datasets generated by MOOC activity provide a strong basis for this type of approach.

Kop and her colleagues looked at engagement patterns on the PLENK2010 connectivist MOOC [10]. As this type of MOOC links a network of tools and resources, it is difficult to track all activity, but their conclusion was that 40-60 of the 1,641 learners registered on the course contributed actively on a regular basis, and that this engagement supported positive learning outcomes. The visible participation rate of others was much lower, indicating a consuming behaviour. This division between the Visible Contributors and the Consumers signals that there are different ways of interacting on a MOOC, but this classification is strongly

associated with connectivist pedagogy and is less relevant to a course in which key elements are pre-defined course content and assessment.

In 2013, Kizilcec and his colleagues [9] analysed patterns of engagement and disengagement in three MOOCs on the Coursera platform. Two understandings underpinned this analysis. One was the view that ‘learning is a process of individual knowledge construction’. The second was the view that the principal features of these courses were video lectures and assessments. This analysis found four patterns of engagement with these courses:

- **Completing:** These learners completed the majority of assessments.
- **Auditing:** These learners watched most of the videos but completed assessments infrequently, if at all.
- **Disengaging:** These learners completed assessments at the start of the course, then reduced their engagement.
- **Sampling** These learners explored some course videos.

As these clusters were found consistently across three MOOCs, each targeted at students working at a different educational level, it appeared plausible that they could be applied to other courses and that ‘MOOC designers can apply this simple and scalable categorization to target interventions and develop adaptive course features’ [9].

However, the generalizability of these categories is qualified on the grounds that these classifications ‘would make sense in any MOOC that is based on videos and assessments’ [9]. This suggests that they may only be applicable to courses using pedagogies based on course content and on the use of either formative or summative assessment.

In this paper, we investigate whether the same patterns of engagement are found in MOOCs that employ social constructivist pedagogy, or if other patterns of engagement apply. In social-constructivist MOOCs, knowledge is jointly constructed through conversation. Contributing to or reading discussion comments is therefore an important part of the learning process. In these cases, there are three elements to be taken into account: active engagement with course content, active engagement with course assessment, and active engagement with course discussion.

1.2 Replication

Within the hard sciences, replication of previously published work is routine. Within the social sciences, it is unusual. Recent efforts to attempt to replicate key findings in social psychology have been controversial [1].

To our knowledge, replication work in learning analytics is as yet unknown. As an empirical discipline, we believe that the field of learning analytics is well placed to use such rigorous approaches.

To be clear, we are certainly not attempting to impugn the integrity of Kizilcec and his colleagues: we regard their work as important and interesting, and are seeking to test how robust their findings are in a different context.

2. FUTURELEARN DATASET

In order to investigate patterns of engagement within MOOCs employing social-constructivist pedagogy, we used data provided by FutureLearn. This is a company owned by The Open University that is currently in partnership with 39 universities, and organisations such as the British Library and the British Museum, to deliver free online courses. The company has developed a new

MOOC platform based on scalable web technology and underpinned by a social-constructivist pedagogy [14]. Each teaching element (step) is associated with a free-flowing discussion. These are intended to emulate a ‘water cooler discussion’ about the immediate content. As with conversations around an office water cooler, people come and go, nobody is expected to engage throughout, but there can be a continuous sense of lively interaction. Typically, the discussion associated with any step on a FutureLearn MOOC will attract hundreds or thousands of contributions, with ‘Like’ and ‘Follow’ options providing ways of navigating these.

Each FutureLearn partner institution has access to data related to courses they run. We therefore focused our attention on data from four MOOCs run by our institution, The Open University. The data on platform activity are not directly associated with demographic data (as discussed below), and so the figures for gender balance presented in Table 1 are taken from responses to the start-of-course survey for each course. These four MOOCs ran soon after the launch of FutureLearn, at a point when the majority of marketing was taking place within the UK, so the majority of participants on each course were based in the UK.

The number of participants reported in Table 1 indicates the number of individuals – both educators and learners – who appeared in our dataset because they were active on the platform after the start date of the course. People who were active before the start date were removed from the dataset entirely, because the only people with access at that point were educators or FutureLearn staff. In FutureLearn terms, a Fully Participating learner is one who marked a majority of the course steps complete and also completed all the assessment.

Table 1: Overview of dataset

	MOOC1	MOOC2	MOOC3	MOOC4
Subject area	Physical Sciences	Life Sciences	Arts	Business
M	51%	39%	32%	35%
F	48%	61%	67%	65%
Participants	5,069	3,238	16,118	9,778
Fully Participating	1,548	684	3,616	1,416
Participation Rate	31%	21%	22%	14%

Each of these MOOCs specified that no previous experience of studying the subject was required. However, the Arts MOOC was recommended for learners aged 16+, due to the possibility that discussion content might include material that would be inappropriate for minors. Fully Participating learners on these MOOCs could gain a Statement of Participation for a fee, but no credits were awarded and participants were not required to complete assessments by a certain date. As a result, the ‘finished assessment on time’ and ‘finished assessment late’ classifications were of less relevance than on the Coursera study [9]. Coursera course content focuses on videos. While FutureLearn courses also include videos, they make use of substantial quantities of other learning material, predominantly text with pictures.

The different context of Coursera and FutureLearn affected the raw data accessible to us as analysts. Kizilcec *et al* were able to obtain full log data, assessment scores, demographic data, and survey responses, all at an individual learner level. This enabled them to report some interesting correlations. On early FutureLearn

MOOCs, demographic, survey and activity data were not linked by personal identifiers, so we had access to demographic and survey data only at a summary level for the courses we were researching. We were able to obtain individual-level activity data, which were anonymised before we received them. As well as assigning random IDs, this process partially aggregated the activity data, so we only had access to the date and time of a learner’s first visit to a step, not the date or time of any subsequent visits. Graphical inspection of these data suggested that there were no gross differences in activity patterns compared to what one would expect in more traditional logfile analysis, although this change has had some effects, which we note where they arise.

3. CLUSTERING

3.1 Replicating the method

Our initial approach followed the one applied in the Coursera study as far as possible, in order to investigate whether previous findings could be replicated in a different context. That study adopted a methodology set out in [9], designed to identify a small number of ways in which learners interact with MOOCs.

Kizilcec and his colleagues began by computing a description for individual learners of the way in which they engaged in each assessment period of the course (typically a week), and then applied clustering techniques to find subpopulations in these engagement descriptions. FutureLearn courses are divided into weeks and, although there is no compulsion to work through at the same pace as others, an overview of activity data showed that activity did indeed follow a weekly cycle, spiking after the weekly course email was sent out each Monday (see Figure 1). Underlying the spikes, there was a general fall-off in activity over time, as would be expected in a MOOC [3]. It is possible, however, that our measure of activity (time of first visit to each piece of content) masked later activity because re-visits to content visited earlier were not counted, and so Figure 1 is likely to overstate the fall-off in activity.

We therefore followed the Coursera study and computed a description for each learner based on individual activity in a course week. Their ‘engagement descriptions’ have six or eight elements, because each course ran for six or eight weeks.




Figure 1: Activity on MOOC1. The darker (red) lines show the date of the weekly course email.

As in the Coursera study, students' activity in each week was assigned to one of four categories:

- 'T = on track' if they undertook the assessment on time
- 'B = behind' if they submitted the assessment late
- 'A = auditing' if they engaged with content but not with the assessment
- 'O = out' if they did not participate.

Although FutureLearn students were not required to submit assessed work by a particular date, we retained these classifications and counted an assessment as submitted 'late' if it was completed after the end of the course week.

To give an example: if a learner missed the first week of a course (out), submitted assessments on time in weeks 2-6 (on track), submitted week 7's assessment but in week 8 (behind), and then engaged with content only in week 8 (auditing), their engagement profile would appear as [O, T, T, T, T, T, B, A].

Once we had created engagement profiles for each learner in this way, we followed the method of Kizilcec's team and applied the k -means clustering algorithm to partition the learners into a small number of groups. To cluster engagement patterns, a numerical value for the dissimilarity between them is required. There are several possibilities for this, but in this initial exploration, we followed their approach: we assigned numerical values to each label (On Track = 3, Behind = 2, Auditing = 1, Out = 0), calculated the L_1 norm for each engagement pattern, and used that as the basis for 1-dimensional k -means clustering.




Figure 2: Proportion of users on MOOC1 falling into each category, by week.

We also followed their practice in repeating clustering 100 times and selecting the solution with the highest likelihood, because k -means has random aspects.

To attempt to replicate their findings, we focused on extracting four clusters. However, our clusters did not match those found in the Coursera study.

The method produced two clusters that were very similar to the ones they found. We found a cluster strikingly similar to their 'Completing' group: learners who completed almost all the assessments. We also found a cluster similar to their 'Sampling' group: learners who visited only once or twice, and did not attempt any assessment.

However, the other two clusters did not match theirs so well. They found an 'Auditing' cluster 'who did assessments infrequently if at all and engaged instead by watching video lectures', and a 'Disengaging' cluster 'who did assessments at the beginning of the course but then have a marked decrease in engagement'. Our two equivalent clusters did include some learners with patterns like that, but they also included many who did not fit those descriptions neatly, and there seemed to be significant overlap. This was reflected in the silhouette scores: they reported an average silhouette width of 0.8, but our data achieved only 0.67. (The closer to 1.0, the better clustered.)

Selecting a number of clusters, k , to extract using k -means is notoriously problematic, unless there is an unambiguous *a priori* rationale. A notable feature of k -means clustering is that it will always generate k clusters, regardless of whether another number would provide a better fit for the data.

We therefore repeated the analysis using values for k that ranged from 3 to 8. We found that the silhouette width was at a minimum for $k=4$, suggesting that this might be the least suitable number of clusters for our data.

We were concerned that the one-dimensional approach was discarding potentially useful information about patterns of engagement before the clustering algorithm could use them, so we repeated the analysis again, this time running k -means on the numeric engagement profiles directly, treating them as 6- or 8-dimensional vectors. We explored the four-cluster solution in detail, and again found 'Completing' and 'Sampling' clusters, but no clear pattern in the other two.

Again, we explored k from 3 to 8, which yielded silhouette widths rather lower than for the 1-dimensional approach (around 0.4), and which decreased monotonically as k increased.

3.2 Adapting the method

Having had mixed results from direct application of the method, we developed a new classification, in order to reflect the importance of discussion in FutureLearn MOOCs. For each content week, students were assigned the value 1 if they viewed content, 2 if they posted a comment, 4 if they submitted the week's assessment late and 8 if they completed the final assessment before the end of the week in which it was set (i.e. early, or on time). These values were added up to give a total for each week.

The majority of comments on FutureLearn are associated with a content step, as study weeks typically include only one or two discussion steps. There is therefore no way to tell from our log data whether learners had participated in the discussion by reading but not by adding a comment, so this option could not be coded.

The close link between content and comments also meant it was not possible for learners to add a comment without viewing content (or the introduction to a discussion step). In theory this would also make scores of 2, 6 and 10 impossible. However, since we were restricted to 'first visits' data, these scores proved to be possible in practice, as returning to a piece of content to post a comment on it (perhaps a follow-up to a previous point, or a reply to another learner) does not show up in our activity data. These scores were rare, however, and they represent only 1.7% of the scores assigned on MOOC1.

The possible scores that could be assigned to an individual learner in any one week are set out in Table 2 (below).

Table 2: Scoring method

Score	Interpretation
1	only visited content (for example, video, audio, text)
2	commented but visited no new content
3	visited content and commented
4	did the assessment late and did nothing else that week
5	visited content and did the assessment late
6	did the assessment late, commented, but visited no new content
7	visited content, commented, late assessment
8	assessment early or on time, but nothing else that week
9	visited content and completed assessment early / on time
10	assessment early or on time, commented, but visited no new content
11	visited, posted, completed assessment early / on time

We used the k -means algorithm to extract clusters from the engagement profiles directly, as a 6- or 8-dimensional vector for each learner, to allow for the possibility of clustering by time of activity, as well as by total activity. The variation in dimensions here was because three of these MOOCs ran for eight weeks but MOOC2, the life sciences MOOC, ran for six weeks.




Figure 3: Silhouette widths for $k=3$ to 8

Selecting a number of clusters to extract was not straightforward. A wide variety of techniques was employed,¹ almost all of which yielded unhelpful results, recommending either a minimal number of clusters (2) or a maximal number.

However, two methods suggested that seven clusters should be extracted. The mean silhouette width reached a local maximum at $k = 7$ (see Figure 3). We also used a scree plot to visualise the total within-groups sum of squares (the measure of how closely the clusters group together that k -means seeks to minimize). The method in using a scree plot is to identify a ‘kink’ in the plot, which is somewhat subjective. In this case, the within-groups sum of squares dipped down suggestively at $k = 7$ (see Figure 4).

¹ All the techniques listed in this Stack Exchange post: <http://stackoverflow.com/questions/15376075/cluster-analysis-in-r-determine-the-optimal-number-of-clusters/15376462>




Figure 4: Scree plot of the within-groups sum of squares for $k=1$ to 10

3.3 Analysis of clusters

Our cluster descriptions were developed by analysis of learner profiles on MOOC1. Like Kizilcec *et al*, we wanted to ensure that the clusters made sense from an educational point of view, even though these descriptions were developed after the analysis had taken place. After developing the cluster profiles for MOOC1, we examined the other MOOCs. With three exceptions, there was very close agreement indeed between the clusters found.

Two of the MOOCs had substantive structural differences from the others. MOOC2, the life sciences course, ran for six weeks, rather than eight. MOOC3, the arts course, ran for eight weeks, but included only three assessments, rather than one every week. Learners on MOOC3 posted substantially more comments than on the other courses. Despite these differences, five of the seven clusters were found in substantially identical form in all four MOOCs. MOOC2 and MOOC3 between them generated three clusters that did not match those developed on MOOC1, and they are appended to the list as supplementary clusters.

Cluster descriptions are given below, along with notes of where individual MOOCs differed. The clusters that were found across all four MOOCs are indicated with a *, as well as in the text. The proportion of learners in each cluster is shown in Table 3 (below). In all cases, the term ‘average’ refers to the mean.

Cluster I: Samplers*

Learners in this cluster visited, but only briefly. It is essentially identical to Kizilcec *et al*'s ‘Sampling’ cluster. Samplers made up the largest cluster in all four MOOCs, accounting for 37%-39% of learners (56% on MOOC4). They typically visited about 5% of the course, with a few Samplers (11%-24%) visiting only a single step, although only 1% of Samplers on MOOC3 did so. They were active in a very small number of weeks, often including week 1, but not always – 25%-40% joined the course after week 1. Very few Samplers posted comments (6%-15%), and very few submitted any assessment, although the handful that did so typically did this in week 1.

A typical engagement profile for this cluster is:

[1, 0, 0, 0, 0, 0, 0]

This cluster was highly stable across all MOOCs, and across most values of k .

Cluster II: Strong Starters*

These learners completed the first assessment of the course, but then dropped out.

Strong Starters made up 8%-14% of learners. All of them submitted the first assignment, but then their engagement dropped off sharply, with very little activity after that. A little over a third of them posted comments (35%-38%), and those who did so did not post very many (1.7-4.0), except on MOOC3, where 73% of learners posted an average of 13.7 comments.

A typical engagement profile for this cluster is:

[9, 1, 0, 0, 0, 0, 0]

Again, this cluster was highly stable across all MOOCs and most values of *k*.

Table 3: Proportion of learners in each cluster, by MOOC

Cluster	MOOC1	MOOC2	MOOC3	MOOC4
I Samplers	39%	39%	37%	56%
II Strong Starters	11%	14%	8%	10%
III Returners	6%	8%	-	7%
IV Mid-way Dropouts	6%	-	-	7%
V Nearly There	6%	6%	6%	5%
VI Late Completers	8%	7%	0.2%	6%
VII Keen Completers	23%	13%	7%	9%
MOOC2 Sup	-	13%	-	-
MOOC3 Sup 1	-	-	20%	-
MOOC3 Sup 2	-	-	2%	-

Cluster III: Returners

These learners completed the assessment in the first week, returned to do so again in the second week, and then dropped out.

Returners made up 6%-8% of learners, with the exception of those studying MOOC3, where this cluster did not appear. This is almost certainly because there are only three assessments in that course, with at least one week between each, so this pattern of activity was not possible.

Almost all the learners in this cluster (>97%) submitted the assessment for week 1 and in week 2. This does not mean that they all visited in both weeks; some of them did the week 1 assessment late. No Returners explored all the course steps; the average amount of steps visited varied from 23% to 47%. After the first two weeks, there was very little activity indeed.

A typical engagement profile for this cluster is:

[9, 9, 0, 0, 0, 0, 0]

This cluster showed some variability with varying *k*.

Cluster IV: Mid-way Dropouts

These learners completed three or four assessments, but then dropped out about half way through the course.

Mid-way Dropouts made up 6% of learners on MOOC1, and 7% of learners on MOOC4. This cluster did not appear for MOOC2 and MOOC3, because their unusual structure (shorter, fewer assessments) meant that there were not enough assessments other than the final one for them to have this profile.

These learners visited about half of the course (47%, 59%), and roughly half posted comments (38%, 49%), posting 6.3-6.5 comments on average.

A typical engagement profile for this cluster is:

[9, 9, 9, 4, 1, 1, 0, 0]

This cluster, like the previous one, showed some variability with varying *k*.

Cluster V: Nearly There*

These learners consistently completed assessments, but then dropped out just before the end of the course.

Nearly There learners accounted for 5%-6% of learners on all four MOOCs. They typically visited over three-quarters of the course (82%-80%) and submitted assessments consistently (>90%) until week 5, and mostly on time (40%-75%), after which their activity declined steeply, and few completed the final assessment (3%-17%), none on time.

Many of them posted comments (48%-65%), and those who did posted an average of 5.7-8.3 – except for MOOC3, where 80% of learners posted an average of 21.8 comments each.

A typical engagement profile for this cluster is:

[11, 11, 9, 11, 9, 9, 8, 0]

This cluster appeared definitively for all four MOOCs, but was somewhat variable with varying *k*.

Cluster VI: Late Completers*

This cluster includes learners who completed the final assessment, and submitted most of the other assessments, but were either late or missed some out.

Late Completers accounted for 6%-8% of learners, except on MOOC3, where only 0.2% of learners fell into this cluster. Each week, including the final week, more than 94% of this cluster submitted their assessment. The average proportion submitting late varied from 16% to 59%. However, more than three quarters of them submitted the final assessment on time (78%-90%).

Fewer than half of these learners posted comments (40%-43%), apart from on MOOC3 where 76% did so, and those that posted made an average of 7.9-15.0 comments.

A typical engagement profile for this cluster is:

[5, 5, 5, 5, 5, 9, 9, 9]

This cluster was fairly stable across all MOOCs, and across most values of *k*.

Cluster VII: Keen Completers*

This cluster consists of learners who completed the course diligently, engaging actively throughout.

Keen Completers accounted for 7% to 13% of learners, apart from on MOOC1, where 23% of learners fell into this cluster. All learners in this cluster completed all the assessments, including the final one, and almost all of them on time (>80%). On average, the Keen Completers visited >90% of the course content.

They were also assiduous commenters. About two-thirds of this cluster (68%-73%) contributed 20.8-24.4 comments on average, although again MOOC3 stood out, with all learners commenting, posting an impressive average of 53.7 comments each.

A typical engagement profile for this cluster is:

[11, 11, 9, 9, 11, 11, 9, 9]

This cluster was highly stable across all MOOCs, and across values of *k*.

MOOC2 Supplementary Cluster

Because it was shorter than the other MOOCs, the Mid-way Dropouts cluster was not found on this MOOC. Instead, a rather miscellaneous and hard-to-interpret cluster was generated, accounting for 13% of the learners. Sixty-seven per cent of the cluster submitted the first assessment, but late, and many either visited regularly or left comments. This cluster seems to fall between Samplers and Strong Starters.

MOOC3 Supplementary Cluster 1: Samplers-Who-Comment

MOOC3 had only three assessments, so the Returners and Mid-way Dropouts clusters were not found.

The first alternative cluster was a substantial one, accounting for 20% of the learners. It was very similar to the Samplers cluster, in that the learners visited only briefly – however, all these learners left comments. They left an average of 5.3 comments, which is high compared to the other MOOCs, but low for MOOC3.

A typical engagement profile for this cluster is:

[3, 0, 1, 0, 1, 0, 0, 0]

MOOC3 Supplementary Cluster 2:

The second alternative cluster was much smaller, accounting for only 2% of MOOC3. These learners submitted the final assessment on time, but had not engaged consistently across the course. In this cluster, 44% submitted the first assessment, 22% the second. They visited on average 44% of the course. About half of these learners left comments (53%), at a rate of 10.1 each.

A typical engagement profile for this cluster is:

[3, 0, 0, 1, 1, 0, 0, 9]

4. DISCUSSION

Kizilcec and his team found four clusters within their data; learners who were completing, auditing, disengaging and sampling [9]. The teams identified these clusters after structuring their data in a way that placed emphasis on learner interaction with content, assessment and deadlines. As Section 3.1 showed, we found ‘Completing’ and ‘Sampling’ clusters, but the others were not found so clearly within the FutureLearn datasets.

The socio-constructivist pedagogy of FutureLearn incorporates not only content and assessment within a course’s learning design but also discussion. In addition, as credit is not currently available for FutureLearn courses, assessment is typically used formatively (to support learning) rather than summatively (to gauge how much has been learned). When active contribution to discussion was included within our model, seven distinct clusters emerged. This indicates that this approach is not pedagogy-neutral, and that opportunities for joint knowledge construction should be taken into account where they are available.

The ‘Sampling’ cluster identified in 2013 [9] can be related to two clusters in the current study: the Samplers and the Strong Starters. The Samplers cluster includes people who arrived, visited a few pieces of content and then left. However, it also includes many learners who arrived late – in Week 1, in Week 2 or even later in the course, and who then did not sustain their engagement. In many cases, this late arrival may have indicated that they never intended to engage with the entire course. However, it is likely that many found and joined the course after its start, or were unable to take part in its early days. This suggests that the time-bounded nature of the course, with a cohort working through the material together week by week, discouraged participation by late arrivals.

The Strong Starters appear, in the FutureLearn datasets, to be distinct from the Samplers. On average, they completed the first week, they made one or two comments, and many of them looked beyond the first week’s materials. These learners did more than simply sample content, but were not able to sustain their engagement with the course.

The clusters Strong Starters, Returners, Mid-way Dropouts and Nearly There are chiefly distinguished by the time period over which their members engaged with the course. Strong Starters engaged for a week, Returners for a couple of weeks, Mid-way Dropouts for about half the course, and Nearly There for all but the last week or two. Not surprisingly, an increase in time spent on the course was associated with a rise in the mean proportion of the course visited by group members and a rise in the mean number of comments contributed.

The two clusters Late Completers and Keen Completers can be considered as a pair, because they include the learners who have engaged with the majority of the material and all the assessments and who are therefore classified by FutureLearn as Fully Participating learners. The main characteristic that differentiates the two clusters is the number of comments posted. All the clusters considered previously in the analysis averaged about one post per person per week of engagement. On average, Samplers engaged for less than a week and posted less than one comment. Strong Starters participated for just over a week and posted just over one comment. Returners participated for a little over two weeks and posted just over two comments. Learners in the Nearly There cluster engaged for four to six weeks and posted around five comments. The Late Completers cluster also broadly followed this one-comment-per-week average. However, the Keen Completers posted twice as much as those clusters, averaging well over two comments per week.

It seems, then, that there were two dominant approaches to completing the course. The more popular approach (there were consistently more Keen Completers than Late Completers) was to engage fully. Keen Completers did almost everything, every week, and visited almost every step of the course. The Late Completers group did not engage with every aspect of the course: they commented less, and they were more frequently late in submitted assessments. So engaging actively with the comments is associated with a more extensive engagement with the course materials and assessment than is apparent in other clusters – although further work is required to explore the causal relationships at work.

5. CONCLUSION

5.1 Implications for practice

The key characteristic of learning analytics, as distinct from the general category of quantitative educational research, is a direct concern with improving learning and learning environments. ‘The key step is “closing the loop” by feeding back [...] to learners through one or more interventions’ [2]. How can this analysis achieve that?

The clusters identified here can help inform a range of strategies for intervention and improvement.

Providing previews of course material would allow Samplers to make a more informed decision about whether to register in the first place.

The four clusters in which learners engage solidly but fail to complete – the Strong Starters, Returners, Mid-way Dropouts, and

particularly the Nearly There cluster – are precisely those learners that educators would want to focus on to improve completion. This data analysis suggests that there are several points at which learners may leave a course, and these may be dealt with in different ways.

Discussion steps set up for latecomers could be used to support those who fall behind at the start, while sign-up pages could draw attention to the difficulties encountered by learners who move through the course out of step with the cohort. These pages might suggest that potential learners return later and register for a subsequent presentation. Free-text survey comments on FutureLearn MOOCs that have run more than once show that some learners return to subsequent presentations in order to finish working through the material with a group – even though they already have access to the same content because they registered on the original presentation.

The course week is an artificial distinction that imposes a weekly format on the learning design of the course. In doing so, it makes the decision to move on to the next week of the course appear more significant than the decision to move on to the next step. Changes to learning design could provide bridges between course weeks, stressing links between these weeks and pointing learners forward. This might reduce the drop off that is experienced by these courses, particularly at the ends of Week 1 and Week 2.

The two most clearly successful clusters – the Late Completers and the Keen Completers also afford actionable insights. In the context of FutureLearn, at least, they appear to provide some degree of vindication of the extensive and structured use of discussion comments.

5.2 Implications for future research

Two broad clusters of learners appear to be highly robust across the number of clusters extracted and across the different contexts: ‘Sampling’/‘Samplers’, and ‘Completing’/‘Keen Completers’. However, this work suggests strongly that it is not possible simply to take a clustering approach from one learning context and apply it in another, even when the contexts are apparently similar. Convincing clusters of learners are most likely to emerge when the approach is informed by the texture of the learning context.

Our experience here leads us to suspect that it is relatively easy to produce good ‘storytelling’ from data clusters, with relatively little dependence on the quality of fit of those clusters. We shall certainly be cautious in our claims in future.

This paper has explored a single clustering method: *k*-means. There are many others, and, given the difficulty identified in robustly deciding the number of clusters to extract, it seems likely that a hierarchical clustering approach might prove more suitable.

In addition, the clustering here has used only the default *k*-means metric for assessing similarity (Euclidean distance). Assigning numeric values to the learner engagement profiles was required to make this possible. Future work could explore clustering on textual profiles directly, using Levenshtein distance.

5.3 Final remarks

As open offerings, for some values of ‘open’, MOOCs afford a wide range of engagement patterns. Two have emerged clearly from this work: the sampling behaviour employed by people who visit the course briefly, and the completing behaviour employed by learners who complete the course thoroughly. The patterns in between those two extremes are likely to prove the richest for

improving the quality of learning and the learning environment, but are also, at this point, the least clear and compelling.

6. ACKNOWLEDGEMENTS

We would like to thank the FutureLearn team for kindly giving us access to the activity data that were used in this analysis.

7. REFERENCES

- [1] BOHANNON, J., 2014 (23 May). Replication effort provokes praise – and ‘bullying’ charges. *Science*, 344, 6186.
- [2] CLOW, D., 2012. The learning analytics cycle: closing the loop effectively. In *LAK12, 29 April - 2 May* Vancouver, BC, 134.
- [3] CLOW, D., 2013. MOOCs and the funnel of participation. In *LAK 2013, 8-12 April* Leuven, Belgium, 185-189.
- [4] DANIEL, J., 2012. Making sense of MOOCs: Musings in a maze of myth, paradox and possibility In *Journal of Interactive Media in Education*, 18. Retrieved from <http://www-jime.open.ac.uk/jime/article/view/2012-18>
- [5] DOWNES, S., 2012. Connectivism and Connective Knowledge. Retrieved from http://www.downes.ca/files/books/Connective_Knowled ge-19May2012.pdf
- [6] DOWNES, S., 2014. Like Reading a Newspaper (21 March). In *Half an Hour*. Retrieved from <http://halfanhour.blogspot.co.uk/2014/03/like-reading-newspaper.html>
- [7] JORDAN, K., 2014. Initial trends in enrolment and completion of massive open online courses. *International Review of Research in Open and Distance Learning* 15, 1, 133-160.
- [8] JORDAN, K., 2014. MOOC Completion Rates: The Data. From www.katyjordan.com/MOOCproject.html
- [9] KIZILCEC, R.F., PIECH, C., and SCHNEIDER, E., 2013. Deconstructing disengagement: analyzing learner subpopulations in massive open online courses. In *LAK 2013 (8-12 April)*, Leuven, Belgium, 170-179.
- [10] KOP, R., FOURNIER, H., and SUI FAI, J.M., 2011. A pedagogy of abundance or a pedagogy to support human beings? Participant support on Massive Open Online Courses. *The International Review of Research in Open and Distance Learning* 12, 7, 74-93.
- [11] LAURILLARD, D., 2002. *Rethinking University Teaching: a Framework for the Effective Use of Learning Technologies*. RoutledgeFalmer, London.
- [12] LITTLEJOHN, A., 2013. *Understanding Massive Open Online Courses*. CEMCA.
- [13] PASK, G., 1976. *Conversation Theory: Applications in Education and Epistemology*. Elsevier, Amsterdam and New York.
- [14] SHARPLES, M. and FERGUSON, R., 2014. Innovative Pedagogy at Massive Scale: Teaching and Learning in MOOCs. In *EC-TEL 2014 (16-19 September)* Lecture Notes in Computer Science, vol. 8719, 98-111.
- [15] SIEMENS, G., 2012. MOOCs are Really a Platform. www.elearnspace.org/blog/2012/07/25/moocs-are-really-a-platform.
- [16] SOLAR, 2011. Open Learning Analytics: An Integrated & Modularized Platform. White Paper. In *Society for Learning Analytics Research*. Retrieved from <http://solaresearch.org/OpenLearningAnalytics.pdf>