Can Learner Characteristics Predict their Behaviour on MOOCs?

1st Author	2nd Author	3rd Author
1st author's affiliation	2nd author's affiliation	3rd author's affiliation
1st line of address	1st line of address	1st line of address
2nd line of address	2nd line of address	2nd line of address
Telephone number, incl. country code	Telephone number, incl. country code	Telephone number, incl. country code
1st author's E-mail address	2nd E-mail	3rd E-mail

ABSTRACT

Stereotyping is the first type of adaptation in education ever proposed. However, the early systems have never dealt with the numbers of learners that current MOOCs provide. Thus, the umbrella question that this work tackles is if learner characteristics can predict their overall, but also fine-grain behaviour. Earlier results point at differences related to gender or to age. However, our finer-grain analysis shows that the result may further depend on the course topic, or even week. Surprisingly, for instance, women chat less in a Psychology-related course, but more (or similar) on a Computer Science course. These results are analysed in this paper in details, including different methods of averaging comments, leading to surprisingly different results. The outcomes can help in informing future runs, in terms of potential personalised feedback for teachers and students.

Keywords

Learner characteristics, stereotypes, MOOCs, FutureLearn, online behaviour prediction

1. INTRODUCTION

Stereotyping is one of the earliest user modelling approach to adaptation and recommendation. It was first introduced by Rich in a book recommender system, Grundy [1], which builds models for individual users based on personal information gathered through interactive dialogues. A stereotype is a collection of physical characteristics or frequently occurring characteristics of individual users such as gender, age, engagement, performance and so on. Creating stereotypes has become a common approach to user modelling – it uses a small amount of initial information to assume a large number of default assumptions [2] which may be updated when more information about individuals becomes available [1].

Stereotyping has been criticised as being too simplistic, and then, again, applied, due to its simplicity. With the advent of the MOOCs, past stereotypes can be evaluated once again at a much larger scale than by preceding research, and confirmed or infirmed. Whilst MOOCs have started being analysed more thoroughly in the literature, few researches, as will be seen, are looking into the temporal, fine-grained analysis of the behaviour, and establishing any relation between the learner behaviour and learner stereotypes.

Our main purpose with this research is to predict the learner overall and fine-grain behaviour based on learner characteristics. In this paper, we specifically focus on the gender stereotype, and its relation to the way learners comment in a MOOC. We base our study on a truly massive FutureLearn course collection of 7 courses delivered via 27 runs between 2012-2016.

The remainder of this paper is organised as follows. We first discuss related research in Section 2. Then, in Section 3, we present the methodology applied, including proposing taking a closer look at what is computed when average comments are measured. Section 4 presents the results, and Section 5 contains the discussion and conclusions.

2. Related Research

As in educational systems, there are two types of stereotyping: fixed and default [2]. A fixed stereotyping classifies learner based on their performance into predefined stereotypes which are determined by, for example, their academic level. In a default stereotyping, a learner is usually stereotyped to default values at the beginning of a learning session; then the settings of the initial stereotype will be gradually altered as the learning process proceeds and more behavioural data is collected [3].

A large body of research has been conducted to explore whether and how learner characteristics can predict their behaviours. Jeske et al. [4] suggest that self-reported learning characteristics can add an important perspective on why and how different learners have different patterns of performance and behaviour while learning. Packham et al. [5] find that successful learners are female, aged between 31 and 50, regardless of their educational level and employment status. Ke and Kwak [6] report that older learners invest more time in online participation. González-Gómez et al. [7] suggest that males have more positive attitudes towards online learning due to their higher computer self-efficiency. Many earlier results point at differences of behaviours related to characteristics such as age and gender.

Over the last six years, massive open online courses (MOOCs) have become increasingly popular and their scale and availability enable a diverse set of learners worldwide to take online courses. In the meanwhile, the amount of learner data collected, including demographic data and behavioural data, has also been increasing. This provides an unprecedented opportunity to further explore the influence of learner characteristics on their behaviours. One approach to understanding learners on MOOCs is by identifying groups of learners with similar behavioural patterns [8] such as clustering learners using engagement factors including the number quizzes attempted [9], [10],. Chua et al. [11] and Tubman et al.

[12] analyse learner commenting behaviours to explore patterns of discussion that occur in MOOCs.

On the other hand, comments have been studied in many setups, including MOOCs. [13] emphasises the importance of using machine learning methods to analyse MOOCs comments, in order to detect the emotions of learners and predict the popularity of each course. [14] focused on grouping students based on their preferences by conducting an online pre-course survey. According to these groups, the relationship between gender showed that females preferred asynchronous text-based posts more. [15] investigated the dropout rate via analysing two MOOC courses with 176 learner's comments on different objects (video, articles, exercises etc.).The study indicated that learners with no negative comments are likely to drop the course very soon. [15] explored the relationship between sentiment ratio measured based on daily forum posts and the number of learners who dropout each day. The study recommended to use sentiment analysis with caution while analysing noisy and quantity-limited comments.

Our study examines how basic learner characteristics, such as gender, can influence learning behaviours such as the patterns of making comments.

3. Methodology

3.1 Terminology

FutureLearn is a MOOC online education platform that provides courses upon weekly basis. Each weekly learning unit consists of several steps, which can be an article, discussion, video or a quiz. The website also allows learners to comment on any given step.

3.2 Data Collection

When a learner joins FutureLearn for the first time, they are directly prompted to complete a survey about their characteristics. Existing learners are also prompted to complete this data, if missing. All the question on the survey are optional and they aim to extract certain information about a learner's gender, age group and education level. In parallel, the system generates logs " to correlate unique IDs and time stamps to learners" recording learner activities such as steps visited, completed, comments added or questions attempts.

3.3 Dataset

The current study is analysing data extracted from 27 runs of 7 MOOCs courses, of 4 main topics: literature (Literature and Mental Health (LT): 6 Weeks), Shakespeare and his world (SP): 10 Weeks; psychology (The mind is flat (MF): 6 Weeks), Babies in mind (BIM): 4 Weeks; computer science (Big Data (BD): 9 Weeks), , and business (Leadership (LS): 6 weeks and Supply chains (SC): 6 Weeks) delivered through FutureLearn by the University of [name-removed]. The study covers 19425 female and 6648 male enrolled learners, out of which 11473 female and 3802 male learners have accessed the course material at least once, and out of which 6240 females and 1833 males have commented at least ones. The material overall has a total number of 2590 steps.

3.4 Formulas

This paper focuses on comments of female and male learners. In order to obtain fine-grain, temporal results, we have analysed comments on a weekly basis. We have also looked at raw numbers; however, to compare on a fairer basis, we have averaged the comments of males and females, computed via the four versions of formulas below.

Version 1: global average (NFE/ NME). Computing behavioural activity based on the global number of students (female/ male) that enrolled in the course.

$$NFE(wi) = (\frac{1}{N^F}) \sum_{k=0}^{N^F(wi)} NComm^F\,(wi) \tag{1a}$$

Where N^F is the total (global) number of females enrolled in the course, for all runs; the rest of the parameters is as defined above. $NComm^F$ is the number of comments posted by females; f(wi) refers to a function f applied to week i. For males, the formula is:

$$NME(wi) = \left(\frac{1}{N^M}\right) \sum_{k=0}^{N^M(wi)} NComm^M\,(wi) \tag{1b}$$

Where N^M is the total (global) number of males enrolled in the course, for all runs; the rest of the variables is the same as above. Whilst these formulas seem the most obvious ones, they include all students who enrol and never actually access the course. To alleviate this, the next version is proposed.

Version 2: access average (NFA/NMA). Computing behavioural activity based on the global number of students (female/ male) active in the course by accessing it. For females, the average is:

$$NFA(wi) = \left(\frac{1}{NA^F}\right) \sum_{k=0}^{N^F(wi)} NComm^F\,(wi) \tag{2a}$$

Where NA^F is the total (global) number of females that have accessed the course, for all runs; the rest of the parameters is as defined above. For males, the average is:

$$NMA(wi) = \left(\frac{1}{NA^M}\right) \sum_{k=0}^{N^M(wi)} NComm^M\,(wi) \tag{2b}$$

Where NA^M is the total number of males who have accessed the course, for all runs; the rest of the parameters is as defined above. Whilst this formula may alleviate some issues, the numbers still include many students who may access the course, but have never commented on it. As the goal here is to analyse comments in particular, the next formula deals with this issue.

Version 3: commenting average (NFC/NMC). Computing behavioural activity based on the global number of students (female / male) active in the course by commenting (at some point – not necessarily that week) in it. For females, the average is:

$$NFC(wi) = \left(\frac{1}{NC^F}\right) \sum_{k=0}^{N^F(wi)} NComm^F\,(wi)$$

(3a)

Where NC^F is the total (global) number of females that have commented the course, for all runs, at some point; the rest of the parameters is as defined above. For males, the average is:

$$NMC(wi) = \left(\frac{1}{NC^M}\right) \sum_{k=0}^{N^M(wi)} NComm^M\,(wi)$$

(3b)

Where NC^M is the total (global) number of males that have commented the course, for all runs, at some point; the rest of the parameters is as defined above. Learners in MOOCs don't behave as students in regular courses: they could comment one week, and not the other, and may be very inconsistent in that. In order to catch also these variations, the next version of the formula is proposed.

Version 4: weekly average (NFCW/NMCW). Computing behavioural activity based on the weekly number of students (female/ male) active in the course for that week, in the sense of those commenting that week. For females, the average is:

$$NFCW(wi) = \frac{\sum_{k=0}^{N^F(wi)} NComm^F\,(wi)}{N^F(wi)}$$

(4a)

Where $N^F(wi)$ is the number of females, for all runs, commenting in week wi; For males, the formula becomes:

$$NMCW(wi) = \frac{\sum_{k=0}^{N^M(wi)} NComm^M\,(wi)}{N^M(wi)}$$

(4b)

With $NComm^M$ the number of comments posted by males and $N^F$the number of males; the rest of the parameters is as defined above.

Comparing the Versions: As can be seen, with the exception of Version 4, where we divide by a changing variable, for the rest of the versions, we divide via a constant, so the shape of the resulting graphs would be the same (although the overall values would change, depending on the number of women enrolled/ accessed/ or having commented in general on the course). As the number of students enrolled is greater than the number of students who access the course, as a great proportion of enrolled students often never access that course; and, respectively, this is further greater than the number of students who comment (some students just 'lurking' in the background, without committing); finally, the latter is greater than the commenters for each week - for females and males, respectively - we have:

$$N^F > NA^F > NC^F > N^F(wi)$$

(5a)

$$N^M > NA^M > NC^M > N^M(wi)$$

(5b)

Thus, the following inequations also hold:

$$NFE(wi) < NFA(wi) < NFC(wi) < NFCW(wi)$$

(5c)

$$NME(wi) < NMA(wi) < NMC(wi) < NFMCW(wi)$$

(5d)

4. Results

4.1 Overall Comments per Gender

Figure 1 shows the numbers of students who were enrolled on average on each course. The most popular courses were clearly on the literature topic. However, of the 6099 students enrolled on the LT course over its 3 runs, only 4214 (69%) students accessed the course at all. Furthermore, only 2513 of those students made any comments. Furthermore, although the Psychology course MF was one of the most popular courses to enrol on, only 26.5% of those enrolled on the course accessed it.

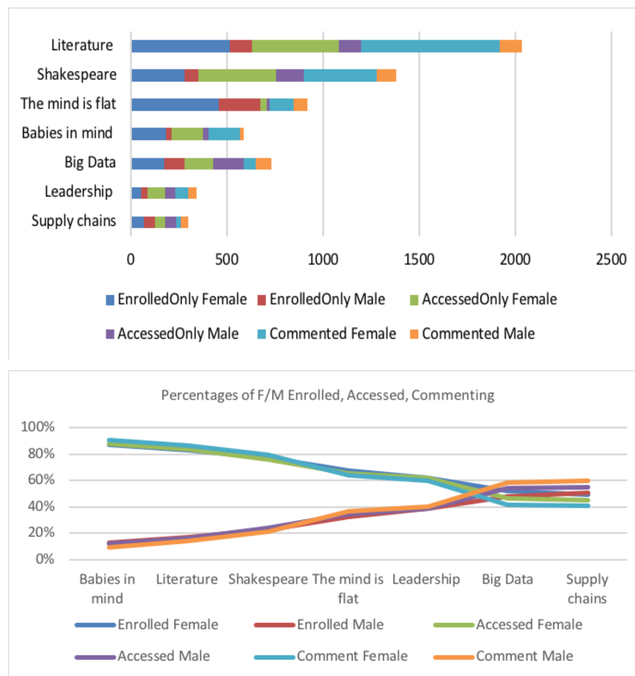

Figure. 1. Number (up) and Percentage (down) of Female and Male students per run of each course, split by different levels of engagement

The lower side of the figure further shows that, whilst numbers of enrollment, access, commenting varied, the percentages that remained after each step were similar within gender, and quite different, between genders. I.e., about 80% of the female learners in the BIM (psychology) course stayed on an accessed the course. A similar percentage of female learners went on and commented. For the same course, the male learner percentage is at under 20%, for enrollment, but also accessing and commenting. Thus, this figure clearly shows that we expect the greatest differences for the literature courses, and that the differences at the other end would possibly be more blurred.

Figure 2 further demonstrates the attrition between students who enrolled on the courses and those that accessed the course. This data is broken down into gender, so that it can be seen, e.g., that a higher percentage of the male learners accessed the BD (Computer Science) and SC (business) courses than the female learners. However, a higher percentage of the enrolled female learners accessed the BIM (psychology) and LT (literature) than the enrolled male learners.

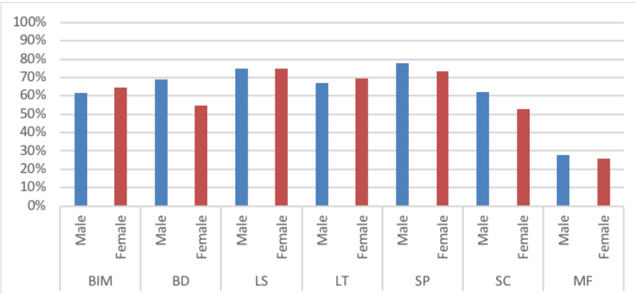

Figure. 2. Percentage of enrolled students who accessed their course

Figure 3 looks at the subset (%) of students who, after accessing a course, left a comment (Version 3, Section 3.4). Although Fig. 2 shows that only around a quarter of MIF (psychology) enrolled learners ever accessed the course, Figure 3 shows that those students who did access the course were much more likely to make a comment about it than students in other courses. The figure also shows that 50.2% of the female learners who accessed the BIM (psychology) course made at least one comment (at some point in the course); however, only 37.77% of the male learners who accessed the course left any comments. Similarly, a higher proportion of female learners on the LT and SP courses (literature) made any comments than the proportion of male learners of those courses. However, in the BD (Computer Science), LS, SC (Business) and the MIF (psychology) courses, proportionally more of the male learners made a comment than the female learners.

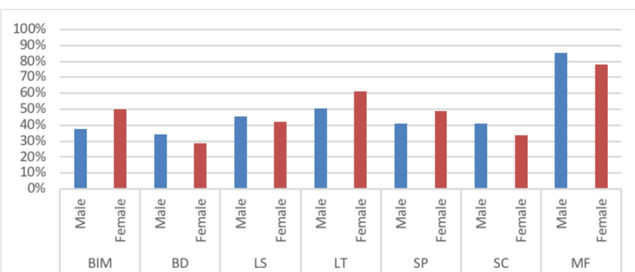

Figure. 3. Percentage of accessing students who made any comments on the course

4.2 Average Comments per Learner

Whereas the above results look at the proportion of male and female learners who made comments, the analysis also looked at how many comments were made for each course, at the fine granularity level of the week. This analysis considers the average number of comments made by all learners who commented on the course at least once (solid line; version 3 in Section 3.4), and all learners who accessed the course at least once (dotted line; Version 2 in Section 3.4); additionally, male learners are shown with a blue line and female learners are represented by a red line.

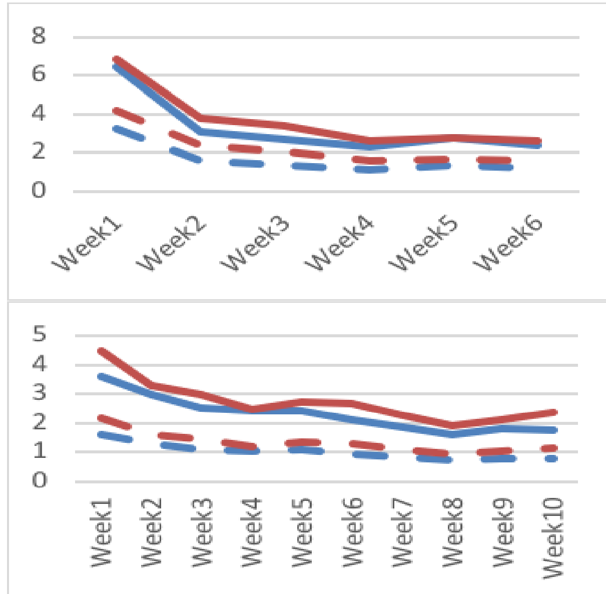

Figure. 4. Literature topic (SP top; LT bottom): comments per learner (version 2 -solid & 3 - dotted; female -red/ male - blue)

Fig. 4 shows that for the SP and LT courses, on average, there were more comments made by female learners than male learners. For Version 2, this difference is consistently statistically significant ($p < 0.05$), but for version 3 the difference is only significant for weeks 2, 3 and 6 (LT) and for weeks 1, 3, 6 and 7 (SP).

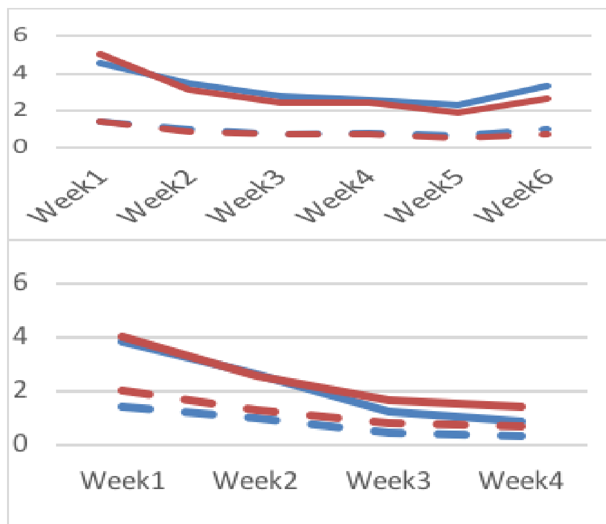

Figure. 5. Psychology topic (MF top; BIM bottom): comments per learner(version 2-solid&3 - dotted; female-red/ male-blue)

Fig. 5 shows a close gender balance for the MF course. However, for weeks 3, 5 and 6 there is a statistically significant ($p < 0.05$ for the Wilcoxon signed rank test) difference when considering only the subgroup of learners who made any comment (version 3). For the BIM course, on average, female learners made more comments than male learners, although not statistically significantly so. However, when considering all learners who accessed the course (Version 2), there is a significant difference for every week ($p < 0.05$ for the Wilcoxon signed rank test).

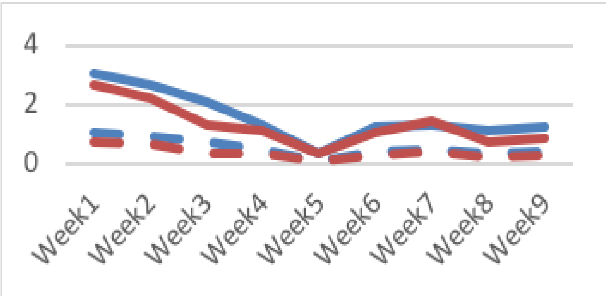

Figure. 6. Computer Science topic (BD): comments per learner (version 2-solid&3 - dotted; female-red/ male-blue)

Fig. 6 shows that male learners of the "Big Data" course made on average more comments than female learners. None of these differences is statistically significant, apart from Week 3 ($p < 0.05$ for the Wilcoxon signed rank test). This significance occurs when considering both subgroups. During week 7, there were more comments made by female learners than male learners, however this is not statistically significant.

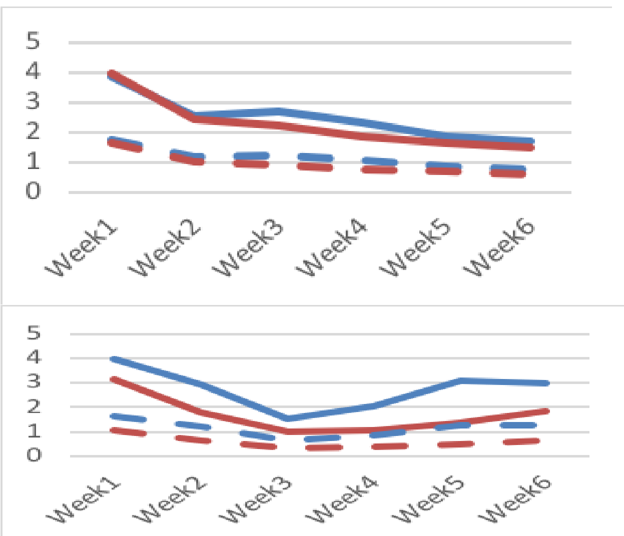

Figure. 7. Business topic (LS top; SC bottom): comments per learner(version 2-solid&3 - dotted; female-red/ male-blue)

Figure 7 shows that male learners of both business courses made on average more comments than female learners, but none of these differences are statistically significant. The only statistical significance ($p < 0.05$) relates to weeks 2 and 6 for SC, when considering Version 2.

4.3 Average Comments per Learner each Week

Whilst the above analysis demonstrates the numbers of comments per learner, as explained in section 3.4, it can also be interesting to analyse how productive the commenting students area in each week (as per Version 4, section 3.4).

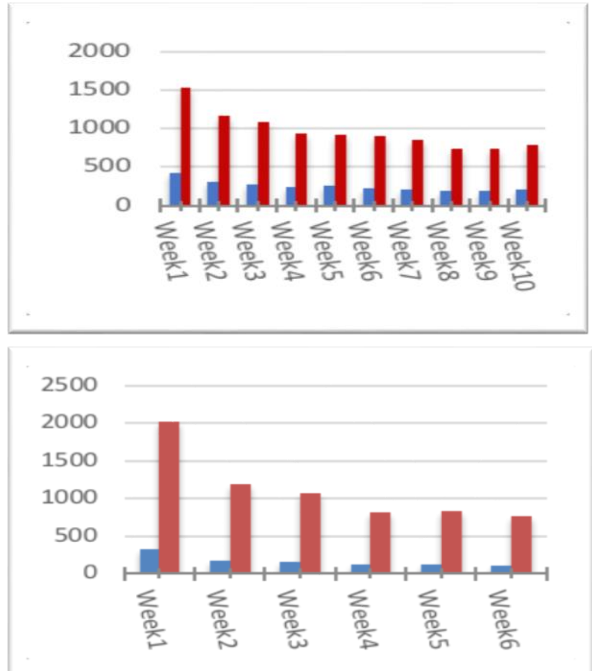

Figure. 8. Literature topic: number of commenters (female - red; male - blue) (SP top; LT bottom)

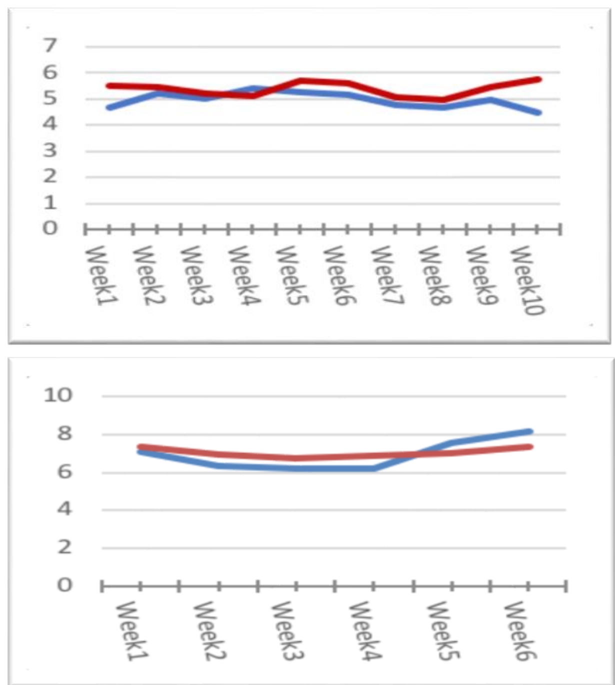

Figure. 9. Literature topic (SP top; LT bottom) comments per learner (version 4)

Figures 8, 9 show that there were consistently many more female commenters than male commenters for literature topic. There was however no difference in the number of comments made by male or female commenters, apart from in week 10, for the SP course, when there was a statistically significant difference ($p < 0.05$), with female commenters on average posting more comments than male commenters.

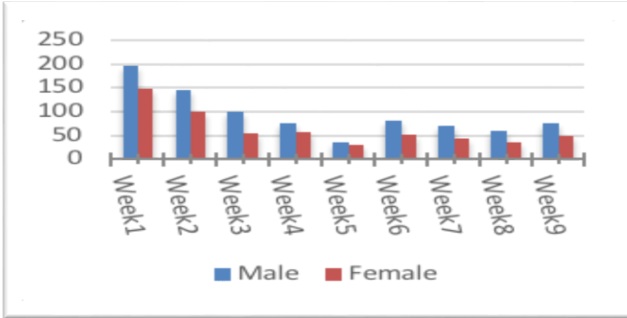

Figure. 10. Computer Science topic (BD) number of commenters (female - red; male - blue).

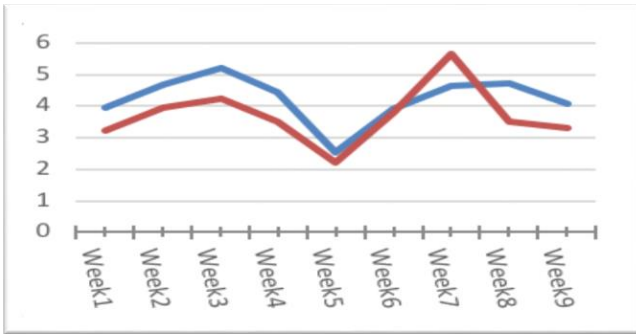

Figure. 11. Number Computer Science topic (BD) comments per learner (version 4).

Figures 10, 11 show that there was a consistently higher number of male learners commenting each week on the Computer Science course. There were slightly fewer comments from each female learner that commented in that particular week apart from in week 7. However, none of these differences is statistically significant.

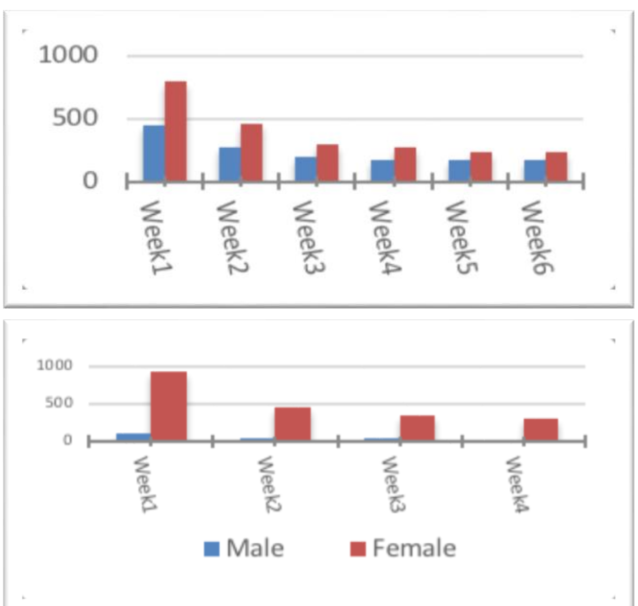

Figure. 12. Psychology topic: number of commenters (female - red; male - blue) (BIM top; MF bottom)

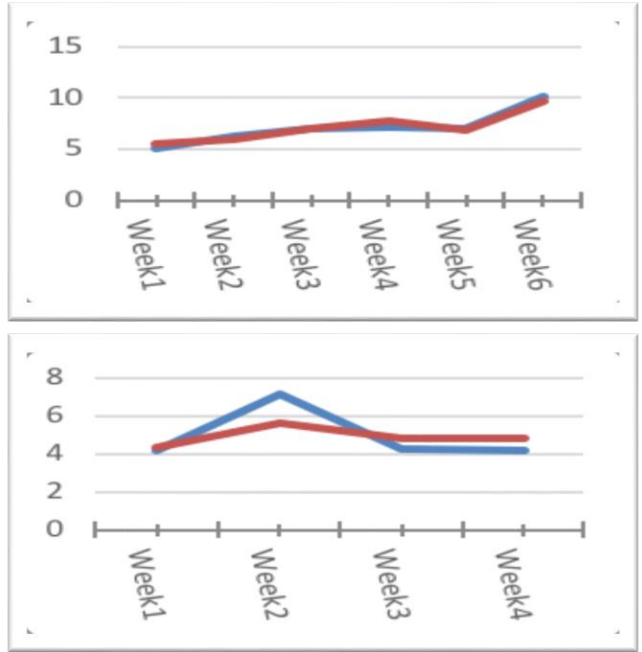

Figure. 13. Psychology topic (BIM top, MF bottom): comments per learner (version 4)

Figures 12, 13 show that there were consistently more female learners commenting than male learners on psychology courses. However, the number of comments made by female and male commenters was not significantly different.

Interestingly, Figure 13 shows that for the BIM course, posting male learners in week 2 (version 4) (of which there were 39 learners), posted more than 7 times, which is more than the average female learner commenter (of which there were 457 learners), posting only 5.65 times. Thus commenting male students were more productive in week 2 only, although this difference is not statistically significant. This result is interesting to compare with Fig. 5, where overall (based on versions 2,3) makes were less productive.

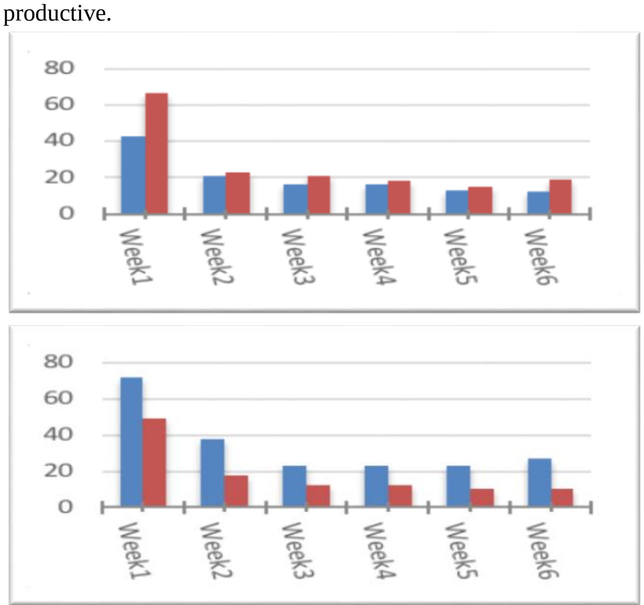

Figure. 14. Business topic: number of commenters (female - red; male - blue) (LS top; SC bottom).

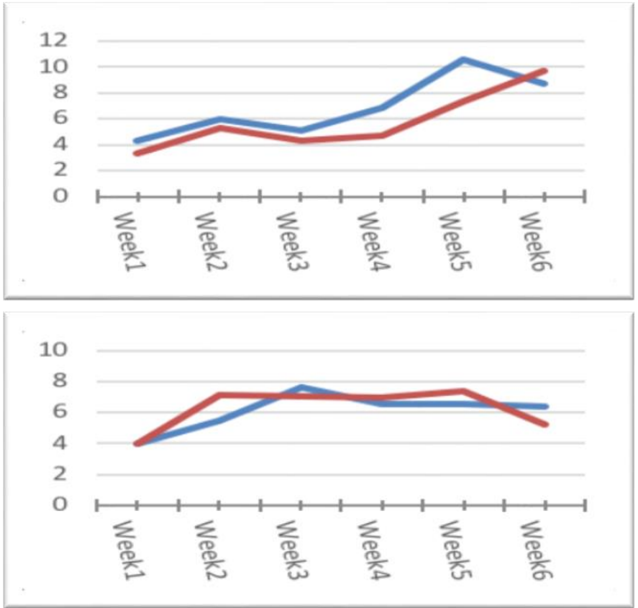

Figure. 15. Business topic (LS top, SC bottom): comments per learner (version 4)

Figures 14, 15 show differences between the business topic courses: for LS, there were more female learners making comments than male learners, whereas for SC, there were more male. However, there was no statistical significance between the number of comments made by male learners and female learners each week.

5. Discussion and Conclusion

The analysis in this paper has highlighted a number of issues which may have been predictable, as well as a few surprises. Firstly, overall, in the courses we have analysed, there are generally speaking more females registered than males. We have also been able to make statements with statistical significance, in general, for the larger courses, such as the literature courses, which were the most popular, followed by the Psychology courses. Computer Science courses are only marginally more popular than Business courses, in our selection.

We have shown that grouping the courses per topic made sense, and that results were relatively similar with such groups – with a notable difference for the business topic, for version 4 (see Figure 15). The latter may be some special case, or this might need to be revisited, e.g., by a teacher of that subject, to check the appropriateness of the classification, and the match between real and desired outcomes.

Importantly, the way the average of comments per learner is computed influences the significance of the results (and, in some cases, the results themselves). Due to the great differences between learners who are enrolled, learners who access the course, or learners who actually comment, in terms of numbers, the conclusions need to clearly vary, when speaking of one cohort or the other. The fewest learners were the ones considered by version 4; they were interesting, in the sense that they were the most active learners. Also, the graphs displayed the greatest variations between the genders. However, differences between genders with respect to that version were not significant. On the other hand, when considering learners who commented at least once, or learners who accessed the course, the significance of the findings grows. In fact, although not depicted here, the greatest significance is achieved when using version 1 – where all enrolled learners are considered. However, as these include a large volume of learners who haven't even seen the course, we didn't analyse them here in further details, as we considered them less relevant to the discussion on commenting students.

Expectations in terms of volume of comments coming from female or male learners clearly vary thus with the topic of the course. Thus, whilst global statements across courses should best be avoided, it is useful to see how students react to a specific course, and then plan for future runs, accordingly. This would help a teacher better understand how to structure the course in a more gender-neutral way, and be enticing to both genders. Furthermore, learners could be notified of options which are targeted to their respective gender. Specific weeks can be analysed when they are triggering behaviour different from the rest of the course – e.g., week 7 in the Computer Science course (see Figure 11), where more female learners comment; or week 6 on the Business topic (SC; Figure 15).

Concluding, we can state that, overall, whilst the participation of females is clearly larger in terms of absolute numbers, in the relatively varied courses we have analysed, in terms of comments produced by the two genders, the topic of the course, the course itself, and often, the week of the course determines which of the genders is commenting more often. Thus, this study clearly shows that it is not enough to study such data on a global scale, and adding up data over several courses with different topics, and over different weeks, may render deceiving results. This study has found several significant differences in the behaviour of female and male learners, in terms of their commenting frequency, at a very fine granularity level: here, at the level of the week of a course. Hence, further studies should look into how the topic and time scale together influence the behaviour of female and male learners for other courses – as possibly other interesting patterns may emerge. Furthermore, here, we only focussed on one stereotype parameter – gender – and one behavioural parameter – commenting. Future research will include a greater variety of such parameters, for extracting a rich picture of how learner characteristics influence learner behaviour in massive online learning environments.

6. REFERENCES

[1] RICH, E., 1979. User modeling via stereotypes. Cognitive science 3, 4, 329-354.

[2] SHUTE, V.J. and ZAPATA-RIVERA, D., 2012. Adaptive educational systems. Adaptive technologies for training and education 7, 27, 1-35.

[3] KAY, J., 2000. Stereotypes, student models and scrutability. In International Conference on Intelligent Tutoring Systems Springer, 19-30.

[4] JESKE, D., STAMOV-ROSSNAGE, C., and BACKHAUS, J., 2014. Learner characteristics predict performance and confidence in e-learning: An analysis of user behaviour and self-evaluation. Journal of Interactive Learning Research 25, 4, 509-529.

[5] PACKHAM, G., JONES, P., MILLER, C., and THOMAS, B., 2004. E-learning and retention: Key factors influencing student withdrawal. Education+ Training 46, 6/7, 335-342.

[6] KE, F. and KWAK, D., 2013. Online learning across ethnicity and age: A study on learning interaction

participation, perception, and learning satisfaction. Computers & education 61, 43-51.

[7] GONZÁLEZ-GÓMEZ, F., GUARDIOLA, J., RODRÍGUEZ, Ó.M., and ALONSO, M.Á.M., 2012. Gender differences in e-learning satisfaction. Computers & education 58, 1, 283-290.

[8] LIU, Z., BROWN, R., LYNCH, C., BARNES, T., BAKER, R.S., BERGNER, Y., and MCNAMARA, D.S., 2016. MOOC Learner Behaviors by Country and Culture; an Exploratory Analysis. EDM 16, 127-134.

[9] FERGUSON, R. and CLOW, D., 2015. Examining engagement: analysing learner subpopulations in massive open online courses (MOOCs). In Proceedings of the Fifth International Conference on Learning Analytics And Knowledge ACM, 51-58.

[10] 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 and knowledge ACM, 170-179.

[11] CHUA, S.M., TAGG, C., SHARPLES, M., and RIENTIES, B., 2017. Discussion Analytics: Identifying Conversations and Social Learners in FutureLearn MOOCs.

[12] TUBMAN, P., OZTOK, M., and BENACHOUR, P., 2016. Being social or social learning: A sociocultural analysis of the FutureLearn MOOC platform. In Advanced Learning Technologies (ICALT), 2016 IEEE 16th International Conference on IEEE, 1-2.

[13] LIU, Z., ZHANG, W., SUN, J., CHENG, H.N., PENG, X., and LIU, S., 2016. Emotion and associated topic detection for course comments in a MOOC platform. In Educational Innovation through Technology (EITT), 2016 International Conference on IEEE, 15-19.

[14] ZHANG, Q., PECK, K.L., HRISTOVA, A., JABLOKOW, K.W., HOFFMAN, V., PARK, E., and BAYECK, R.Y., 2016. Exploring the communication preferences of MOOC learners and the value of preference-based groups: Is grouping enough? Educational Technology Research and Development 64, 4, 809-837.

[15] WEN, M., YANG, D., and ROSE, C., 2014. Sentiment Analysis in MOOC Discussion Forums: What does it tell us? In Educational data mining 2014 Citeseer.