Modeling Micro-Interactions in Self-Regulated Learning: a Data-Driven Methodology

DOI:
10.1016/j.ijhcs.2021.102625
10.1016/j.ijhcs.2021.102625

Document Version
Accepted author manuscript

Citation for published version (APA):
Yu, H., Harper, S., & Vigo, M. (2021). Modeling Micro-Interactions in Self-Regulated Learning: a Data-Driven Methodology. *International Journal of Human Computer Studies*, *151*, Article 102625. https://doi.org/10.1016/j.ijhcs.2021.102625, https://doi.org/10.1016/j.ijhcs.2021.102625

Published in:
International Journal of Human Computer Studies

Modeling Micro-Interactions in Self-Regulated Learning: a Data-Driven Methodology

He Yu, Simon Harper, Markel Vigo

Department of Computer Science, University of Manchester, United Kingdom

Abstract

We explore whether interactive navigational behaviours can be used as a reliable and effective source to measure the progress, achievement, and engagement of a learning process. To do this, we propose a data-driven methodology involving sequential pattern mining and thematic analysis of the low-level navigational interactions. We applied the method on an online learning platform which involved 193 students resulting in six interactive behaviours that are significantly associated with learner achievement including exploration of the first week's materials and exploration of the forum. The value of including these behaviours in predictive models increased their explainability by 10% and accounted for an overall explainability of 82%. Performance evaluations of the models indicate 91-95% accuracy in identifying low-achieving students. Other relevant findings indicate a strong association between the reduction of the behaviours over time and student achievement. This suggests a relationship between student interface learnability and achievement: achievers become more efficient at using the functionalities of an online learning platform. These findings can provide context to learning progress and theoretical foundations of interventions against unhelpful learning behaviours.

Keywords: Interactive behavior, Interaction pattern, Self-regulated online learning

Email addresses: `he.yu@manchester.ac.uk` (He Yu), `simon.harper@manchester.ac.uk` (Simon Harper), `markel.vigo@manchester.ac.uk` (Markel Vigo)

1. Introduction

While online learning platforms provide students with the convenience to access learning materials at any time and location, concerns relating to the self-regulated nature of these platforms are growing. Tracking engagement, achievement, and abandonment features that occur during the learning process has become a rising research topic in the area of learning analytics. The reasons for this are twofold: first, we can provide context for assessment results and differences in student performance by understanding the differences in their level of engagement [35]. This information is valuable for educators as it can be used to identify students who are struggling at an early stage, with previous research suggesting, using performance predictors, that 50% of students who were close to fail or failed could have been identified before the start of the course [83]; second, for those online learning platforms such as connectivist MOOCs (cMOOCs), the absence of an assessment method is one of the challenges of monitoring learning progress. Monitoring these features may provide an assessment-free alternative.

Previous research has associated interactive behaviour with features of learning on Learning Management Systems (LMS) and online search [89, 58], where it was shown that some student interactions are indicators of engagement and can predict knowledge gain. For instance, the number of learning resources visited have been linked to learner achievement [16], the duration of learning sessions can be used to indicate knowledge status [46], properties of search queries such as length or complexity have been linked to learning outcome [84], and the number of resources visited was shown to positively correlate with knowledge gain [27].

User interactions can be classified depending on the level of abstraction, which ranges from low-level physical events (e.g. key press) to high-level task related events (e.g. completing an assignment) [44]. Recent works demand that online learning platforms such as Coursera or Blackboard generate lower level granularity data with rich semantics to allow researchers to extract more interpretable information about what students do in order to create interventions [56, 32]. Established works use coarse-grained features that combine assessment data and user activity [50, 31] suggesting the validity of behavioural data to measure features of learning. Crues *et al.* (2018) [19] provide a comprehensive overview on these features, which mostly include interactions at the same abstraction level, typically coarse grained events such as the number of posts in the forum. Additionally, it has been suggested that both low-level and high-level interactions should be amalgamated as context may be spread across multiple levels and the composition of these events needs to be taken into consideration to give appropriate interpretations [44]. In the context of online learning, it has also been suggested considering interaction patterns, as the order of interactions can encode valuable information about learning behaviour [17]. Low-level interaction patterns have been described elsewhere as micro-interactions [14], where repeated scrolling, re-visitation, and the number of times a response to an opinion question is changed were helpful in identifying issues such as disengagement, low self-efficacy, and confusion.

Inspired by the above works, we explore these connections in an online learning environment where the lack of formal assessment and the analysis of lower level interactions call for a data-driven approach. Specifically, we investigate whether micro-interactions (i.e. students' interactive behaviours) can be used to measure learning progress, achievement, and engagement. In the context of this work, we define engagement as the student's level of participation in course materials and activities on the learning platform [35], and the number of units completed and whether students received a badge are used to indicate learning progress and achievement respectively. We address the following research questions (RQ):

- **RQ1**. Can we identify micro-interactions that are relevant for online learning?
- **RQ2**. Which micro-interactions are associated with learning progress and achievement?
- **RQ3**. What is the added value of using micro-interactions?

To explore whether interface interactions are associated with learning progress and achievement, higher level task-related behaviours (e.g. watching a video) were generated from UI events (e.g. mouse clicks) by applying data mining techniques and qualitative analysis. Then, we explored the relationship between interactive behaviours and learning progress and achievement. Identifying these relationships provides both additional validation and insights on assessing learning. The contributions of this work are:

- We developed a data-driven methodology to isolate micro-interactions that are indicators of learning progress and achievement. We then showed the feasibility of our approach by applying this method on an online learning platform.

- Our approach identified six interactive behaviours which can explain 82% of the student achievement variation and provide a 91-95% accuracy in recognising low-achieving students. The added value of non-engagement metrics added a 10% of explainability of student achievement.

- In the learning platform that serves as a case study of the methodology, we discovered that interactive behaviours involving exploration of the forum, and exploration of the first week materials before leaving the page are significant predictors of student achievement.

- By discovering negative correlations between the frequency of interactive behaviours and achievement (i.e. the less frequent, the higher the achievement rate), we were able to associate the learnability of the online learning platform with the learning outcome.

2. Related Work

This work builds on the metrics of online learning process (e.g. engagement, achievement), interactive behaviours to predict the features of learning, and sequential pattern mining to identify interactive behaviours.

2.1. Online Learning Metrics

In the context of self-regulated online learning, the definition and measurement of achievement can be a single (or a combination) of test scores and grades [38, 16], course/personal learning goals [30, 86], and course completion [72]. In a recent study, student achievement was measured as personalised learning objectives, and the rates of receiving a certificate significantly improved for students whose learning objectives include receiving a certificate when compared to the entire course population [74].

Engagement has been established to be closely related to learning [51, 69, 90] in that the more students engage with activities the better they perform [50], while low-achieving students show lower level of engagement [24]. It has been suggested that low-achieving students receive fewer gamified badges and are less engaged than other students [24]. It has also been shown that evaluating engagement is key to assessing student retention, learning progress, and test performance [11, 19, 71, 78]. Engagement is often measured by learner activities, and the amount of participation in online learning forums and the interaction with lecture videos are two of the most common measures of engagement. Motivation has shown to be a significant predictor of engagement [12], and it was found that motivation can predict behavioural, emotional, and cognitive engagement [79]. Previous work suggests that awarding badges has the potential to increase student motivation as it provides clarity of the learning goals and guidance on how to reach them [43, 1].

Approximately only 10% of registered students complete an online course [54]. Hence, another frequently explored topic is student dropout/abandonment. Dropouts in online courses refer to the students who discontinue their participation in the course, while abandonment refers to passive participants — students who do not unregister and continue the course without being active participants [73]. Many studies have investigated the reasons students dropout, the most common reasons being lack of intention to complete, personal circumstances, bad course design, limited digital skills, unmet expectations, and negative prior experiences [62, 37].

2.2. Modeling Student Interactions

Interactive behaviours and patterns refer to how the learners interact with an online learning platform and how these behaviours are exhibited over time. These factors have been the focus of previous works which involve assessment and prediction. Web navigation behaviours, for example, have been used to predict the next page students are likely to explore [65], their level of engagement/participation [71], the likelihood to dropout [85], and their level of competence [21, 48, 29]. Similarly, student interactions have been used to evaluate participation behaviours in online collaborative learning [20], as well as to predict gaming behaviours in an intelligent tutoring system [59]. Further to this, informed by learning design, a recent study analysed interactive patterns to better understand student behaviours such as resource transition and review [77], and interaction patterns together with self-regulated learning strategies and self-reported variables were found to be accurate predictors of learner types [55].

It has also been suggested that learning is associated with interactive behaviours such as browsing patterns (e.g. how students transition through different tasks) [36], click patterns [66], exploration strategies (including how students interact with videos) [3], exploration choices [48], and problem-solving strategies [91, 40]. Research has shown that students' grades can be accurately predicted using internal-course interactions combined with student activities and behaviours beyond the learning platform [70]. These works suggest that there are opportunities to relate online interactive behaviours with learning.

2.3. Sequential Pattern Mining

Sequential pattern mining (henceforth SPM) algorithms identify interaction patterns by extracting the sequential combination of interaction data that generates the most frequent patterns for a given *support* (i.e. the percentage of sequences that exhibit the pattern). SPMs have traditionally been used in retail where understanding shopping patterns can help increase profit by improving product allocation and display [5].

SPM algorithms generally produce three types of patterns: frequent, closed, and maximal sequential patterns. Frequent patterns are sub-sequences with frequencies exceeding a specified minimum support. Closed sequential patterns are frequent sequential patterns that are not strictly included in other sequential patterns having the same support. Finally, maximal sequential patterns are frequent sequential patterns that are not strictly included in other frequent sequential patterns. For instance, both pattern A $\{i,j\}$ and pattern B $\{i,j,k\}$ are frequent with their support counts being 0.3 and 0.2, respectively. We assume that for pattern A, its only frequent super pattern is B. In this case, pattern A is a closed pattern as its only super pattern is less frequent than itself. However, pattern A is not a maximal pattern as its super pattern B is also frequent. The pattern B would be a maximal pattern if it has no super patterns which are frequent. For mining sequential patterns, CM-SPADE generally has the best performance [33], while CM-ClaSP and CloSPan are generally better suited for mining closed sequential patterns, and so is VMSP for maximal sequential pattern mining [34].

SPM is frequently used in the area of educational data mining; to identify sequences of events that can distinguish stronger/weaker groups in an online collaboration environment [67], and to extract learning features to classify learner groups [81]. It has been used in recent e-learning studies to improve online learning platforms, investigate online collaboration, and explore self-regulated learning [26, 67, 25]. Previous work has also used SPM to mine student behaviours from length and correctness of steps taken in online programming courses [45]. Compared to other similar purpose data mining techniques such as process mining, SPM is more general as it can be applied to various types of sequences. Both techniques were evaluated qualitatively and quantitatively to predict dropouts in online learning, and it was shown that SPM is better suited to handle the data produced by learning processes for predictive purposes [22].

3. Study: Setting, Apparatus and Data collection

Interaction data was extracted from an online learning platform constructed as a cMOOC (connectivist Massive Open Online Course), a type of MOOC that focuses on participatory learning with an emphasis on collaboration and creation. The platform targeted a specific student group, early career researchers. The courses took place in three waves and ran for four weeks each time: 12 November–16 December 2018, 21 January–17 February 2019, and 17 June–14 July 2019. The learning topic was open science and open research methods. The learning materials were divided into four weeks and three learning modules including 63 micro-learning units. The three modules were:

- Data and information literacy: how to search, evaluate, and manage digital information.
- Communication and collaboration: how to use technologies to interact, share, communicate, and collaborate.
- Content creation: how to develop, integrate, and apply copyright to digital content.

The learning materials included text, videos, illustrations, links to online resources, etc. The platform provided a news page with updates about new modules, a wiki page where students share their reflections collaboratively, and a forum where students discuss the learning materials. Within each page there was a main area containing the learning materials and a left sidebar with links to other pages. For certain learning units the contents also included videos. 193 out of 382 students gave their consent to having their interactions collected. When each wave of the MOOC was finished, an online survey was sent to the students, which had a low response rate (10%). Hence, we should be careful making generalisations from the sample: the majority of students were based in the EU, the gender ratio was balanced and most of them were researchers or PhD students in natural sciences, engineering, and social sciences. We did not associate the demographic data to interaction data.

The platform was instrumented to generate low-level event interactive data, including browser window events such as page loads, and mouse and keyboard interactive data. Some of the events contain additional information, such as mouse coordinates for mouse events. We captured a total number of 281,087 interactive events of the types listed in Table 1 using WevQuery [8]. The low-level interface events generated directly from the platform captured student interactions in detail, which can lead to outputs that contain noise such as unintended/unnecessary student actions [23]. A series of pre-processing steps were conducted to combine or transform such events. For example, events which typically follow each other such as mouse-up and mouse-down, were combined and renamed as mouse-press. Key-down were renamed depending on whether they were commands (i.e. return key) or an alphanumeric key and same consecutive events were also merged. To ensure the interpretability of the event sequences, we grouped them by topic, i.e. the homepage, weekly material page, wiki page, forum, news page, and settings. Each recorded event had a corresponding timestamp, URL of the Web page where the event takes place, and the specific DOM element that triggers the event. Hence, a `event-node-URL` triple represented each event. For example, `mouseinorout+main+wiki+multi` indicates multiple mouse movements in the main area of the wiki page.

We computed two indicators of learning: the badge status and the number of units completed. The instructors of the course awarded an 'Open Science Aficionado' badge to those students that conducted weekly assignments. The specific requirements of the badge were to achieve at least 12 out of 24 attainable points from the following activities: vote on other student's forum post (1 pt.); comment on other student's forum post (2 pt); complete weekly assignment (3 pt.). On the weekly assignment of week 1, students had to read an article about the opportunities and challenges of Open Science. On week 2, students had to use social media to share ideas, ask for advice and find research partners to run a research project. Week 3 was about sharing research using workflows involving DOIs, ORCID, open data repositories such as Zenodo and use of Altmetrics. Finally, on week 4, students had to design a research plan using what they had learned in the previous weeks. While the data of all the participants was used for pattern analysis, we excluded 53 students from the predictive modeling analysis as the first wave of the course did not award any badge, leaving 140 students for the analysis involving badges and 193 for progress analysis.

Table 1: Example of captured student interface events.

Type	Events	Description
Mouse	`mousedown`	Start of mouse click action
	`mouseup`	End of mouse click action
	`mousemove`	Mouse movement
	`mouseover`	Hovering into target
	`mouseout`	Hovering out from target
	`doubleclick`	Double mouse click
	`mousewheel`	Mouse wheel interaction
Selection	`select`	Selection of page content
	`cut`	Content cut
	`copy`	Content copy
	`paste`	Content paste
Keyboard	`keydown`	Start of key press action
	`keyup`	End of key press action
	`keypress`	Key press action
Window	`load`	Page is loaded
	`resize`	Browser window is resized
	`unload`	Window is closed
	`windowfocus`	Browser tab gains focus
	`windowblur`	Browser tab loses focus
	`scroll`	Change of scroll state
Other	`change`	Input element state change
	`contextmenu`	Opening of context menu

4. Methodology

We propose a three-step methodology to mine interaction data, analyse sequential patterns, and generate interactive behaviours of interest. First, we benchmark sequential pattern mining algorithms and different versions of our dataset to identify the combination yielding the most *representative* and *informative* patterns. Once we the set the parameters, we run the algorithms and extract the sequential patterns. Second, we apply thematic analysis to interpret and group the extracted patterns, transforming low-level interactions into higher-level interactive behaviours. Third, the interaction patterns that fall under the emerging themes are used to build user models.

4.1. Benchmarking Algorithms and Datasets

The success criteria for this stage aim at maximising the number of patterns generated and the sequence length. We benchmarked a number of SPM algorithms including CloSpan, CM-ClaSP, CM-SPADE, CM-SPAM, and VMSP, which were provided by the SPMF Java library [34].

To find representative patterns we maximised the number of patterns generated by a higher number of users. Each of the input sequences was constructed with a single student's interaction events across all active sessions. To assess how representative a set of patterns are, we explored the value space of the *minimum support* parameter (i.e. the percentage of students that exhibit a given sequence) in the 0.35–0.8 range as values above 0.5 yielded a small number of patterns and the number of patterns increases dramatically when the support decreases below 0.4. To maximise the amount of information, we maximised the length of the pattern. Longer sequences entail more events, and thus more information. Single-event sequences were therefore excluded.

To explore interaction patterns over time, we split our dataset into sessions. Sessions were split by 5, 30, and 40 minutes of inactivity in between two consecutive events which is in line with what is suggested in the literature [47]. This resulted in three different datasets: `gap-5`, `gap-30`, and `gap-40` respectively. We compared the medians of sequence size of the three datasets with minimum support of 0.35–0.5. The number of patterns and their lengths were plotted for each dataset to identify intersections that maximised the criteria for representativeness and sequence length.

4.2. Thematic Analysis

Thematic analysis is a widely used qualitative analytic method for identifying, analysing, and reporting patterns (themes) within, typically, qualitative data [13]. While the use of thematic analysis has been used to evaluate online learning platforms [57, 63, 6, 64], we employ thematic analysis to systematically find themes on the patterns produced by the SPM algorithms in the previous stage and reduce the high number of sequential patterns (and address **RQ1**), being the first of its kind to adopt this method. We adopted the most common form of thematic analysis, a six-step process described by Braun and Clarke [13]: familiarising with the data, coding where we considered the sequential patterns generated as initial codes, and generating the themes. We conducted an inductive approach to determine themes as the data-driven nature of this research. We first followed the semantic approach

Table 2: Example of theme generation process. The first column gives examples of the initial codes (i.e. sequential patterns). The second column shows the sub-themes generated from initial codes. The third column shows the final themes (i.e. interactive behaviours) generated from the sub-themes.

Codes	**Sub-themes**	**Final themes**
load+week_1 windowfocus+week_1 mouseinorout+main+week_1+multi scrollorwheel+week_1+multi mouseinorout+sidebar_left+week_1+multi video+week_1	load and open the week 1 page multiple mouse movement at week 1 main area multiple scrolling in week 1 page multiple mouse movement at week 1 left sidebar video activity in week 1 page	Explore the week 1 page main area, left sidebar, and interact with videos
windowfocus+week_1 mouseinorout+main+week_1+multi mouseinorout+sidebar_left+week_1+multi video+week_1	open the week 1 page multiple mouse movement at week 1 main area multiple mouse movement at week 1 left sidebar video activity in week 1 page	
load+home windowfocus+home mouseinorout+main+home+multi blurfocus_leavepage+home	load and open homepage multiple mouse movement at homepage main area leave home page	Explore the homepage main area then leave page

to transform each code into a sub-theme as a sentence describing the explicit interaction. After the generation of sub-themes, we followed the latent approach with assumptions about the underlying behaviours and determined the final themes. Each theme was created with a combination and transformation of the sub-themes. Examples of this process can be viewed in Table 2 where the codes are results from the sequential pattern mining (i.e. patterns detected in low-level events) and the final themes are the interactive behaviours student exhibit. The interpretation of the behaviours reveals the strategies used by students such as interacting with videos and forums. An independent coder was involved in the thematic analysis. The inter-rater reliability was particularly high, Cohen's $\kappa = 90\%$. This may be due to the fact that sub-theme and theme analysis was conducted under a set of agreed rules: for example, `load+windowfocus` was defined as 'explore' rather than 'load', 'focus', or 'view', which reduced ambiguity.

4.3. Modeling Student Interaction

To investigate **RQ2**, the behaviours identified through thematic analysis and their corresponding interaction patterns were sought in the original dataset for occurrences of each behaviour within each session for each student. These occurrences were represented in mxn matrices, where m is the number of students and n the number of sessions exhibited by the student. For instance, the matrix for behaviour B can be represented as follows:

$$B = \begin{bmatrix} f_{11} & f_{12} & \dots & f_{1n} \\ f_{21} & f_{22} & \dots & f_{2n} \\ \dots & \dots & f_{ij} & \dots \\ f_{m1} & f_{m2} & \dots & f_{mn} \end{bmatrix}$$

where f_{ij} is the $f$requency or number of occurrences of behaviour B for $student_i$ in $Session_j$. Since not all students have the same number of sessions, $f_{ij} = NA$ on those cells that do not have sessions.

To identify which behaviours were associated with indicators of learning, we computed descriptive statistics of the frequencies including the mean, median, and sum of occurrences. We also considered the sessions where students were 'inactive' (i.e. $f_{ij} = 0$), as dwell time is indirectly related with achievement through engagement [88, 53]. Also, to show the evolution of interactive behaviours, we computed the frequency trend to indicate the general direction in which the number of occurrences is developing. To calculate the trend, we calculated the mean, median, and sum of the occurrences across 3, 5, and 10 surrounding sessions. This is in consideration of the accuracy of calculated trend for a relatively sparse dataset where long periods of inactivity exist. The trend is calculated using two methods: the coefficient of its correlation with the session number variations, and the slope in a polynomial equation with the horizontal line representing the increase of session number [61]. The trend, a combination of the aforementioned correlation and slope, was categorised in six levels of strength ranging from strong negative trend to strong positive trend. In total, we included 46 features for each behaviour as listed in Table 3. We then conducted correlation and regression analysis between these features, and the number of units completed (indicator of progress) and the badge status (indicator of achievement).

5. Results

Benchmarking. None of the five SPM algorithms was found to be problematic in terms of execution time and memory usage. As far as the number of patterns was concerned, for all of the datasets, CloSpan produced the fewest patterns while CM-SPADE produced the most. The ranges of sequence size under the minimum support of 0.5–0.8 was 2–4 for CloSpan and 2–8 for CM-SPADE, CM-ClaSP, and VMSP. In consideration of the next step of the methodology, it was more efficient to select a maximal sequential pattern mining algorithm to eliminate the possible overlapping sequences.

Table 3: Description of the features of interactive behaviours.

Feature	Description
`Average`	The average of occurrences across all sessions
`Sum`	The sum of occurrences across all sessions
`NumGap`	Number of consecutive inactivity sessions
`MaxGap`	Maximum number of consecutive inactivity sessions
`Inactiv`	Number of inactivity sessions in total
`AvInactiv`	Proportion of inactivity sessions in all sessions
`Epi_not0`	Number of active sessions
`Episodes`	Number of online sessions in total
`Average_not0`	Average of occurrences across all active sessions
`Median`	Median of occurrences across all sessions
`T+3/5/10+mean/median/sum`	Trend calculated as coefficient with the mean/median/sum of 3/5/10 surrounding sessions
`Pl+3/5/10+mean/median/sum`	Trend calculated as slope with the mean/median/sum of 3/5/10 surrounding sessions
`S+3/5/10+mean/median/sum`	Trend strength of T3/5/10
`Sp+3/5/10+mean/median/sum`	Trend strength of pl3/5/10

We therefore performed our analysis using the maximal sequential pattern mining algorithm – VMSP on the three datasets.

The dataset with sessions separated by 40 minutes of inactivity (i.e. `gap-40`) is the one that generated the longer sequences exhibited by the largest number of students. This was measured as the point at which minimum support and sequence length are both maximised – which was 0.43 for minimum support. As a result, a total number of 110 patterns were generated with a sequence length range of 2–10.

Thematic Analysis. From the initial 110 patterns, 23 themes emerged from the thematic analysis, as shown in Table 4, including the exploration of the homepage, wiki, and the learning materials of week 1. Video and forum activities, which are known to be associated with learning [18, 10, 15] were involved in six and five of the 23 interactive behaviours, respectively. Exploration of the left bar menu, which contains the links to different resources, is also a recurrent element suggesting intent to explore pages. We also found that 'leaving a page' is present in four of the interactive behaviours, suggesting the action of clicking a link provided in learning materials or the student being distracted from the course (i.e. withdrawing from exploring its

Table 4: Interactive behaviours generated from thematic analysis.

Index	Interactive behaviour
1	Visit the wiki page
2	Explore the wiki page and left sidebar
3	Explore the week 1 page main area, left sidebar, and interact with videos
4	Explore the week 1 page main area, left sidebar, and forum main area
5	Explore the week 1 page main area then leave page
6	Explore the week 1 page main area and interacting with videos
7	Explore the week 1 page main area and left sidebar then leave page
8	Explore the week 1 page main area and left sidebar
9	Explore the week 1 main area and forum main area
10	Explore the week 1 page main area
11	Explore the homepage main area, left sidebar, and interact with videos
12	Explore the homepage main area and interact with videos
13	Explore the homepage main area and left sidebar
14	Explore the homepage main area then leave page
15	Explore the homepage main area and week 1 page main area, left sidebar, and interact with videos
16	Explore the homepage main area then week 1 page main area and interact with videos
17	Explore the homepage main area then week 1 page main area and left sidebar
18	Explore the homepage main area and left sidebar then week 1 page main area and interact with videos
19	Explore the homepage main area and left sidebar then week 1 page main area
20	Explore the homepage main area and left sidebar then forum main area
21	Explore the homepage main area
22	Explore the forum main area and left sidebar
23	Explore the forum main area

contents). Eight interactive behaviours were found to involve interactions on different pages including homepage, materials of week 1, and forum.

Modeling Student Interaction and Analysis. For each of the 23 emergent behaviours, we computed the Spearman and Kendall correlation tests between the features in Table 3 and the two indicators of learning we extracted: the badge status and the number of units completed. We found moderate to strong significant correlations $|\rho, \tau > 0.45|$ between the features of six of the behaviours and the badge status as shown in Table 5. We found no significant correlations between the behaviours and the number of completed units.

From the features in Table 3, `Episodes`[1] (i.e. the number of online sessions in total) correlates positively with whether students received a badge ($\rho = 0.42$, $p < 0.001$), which is in line with literature that suggests the more active a student is, the more likely a badge will be awarded [41, 78, 46]. The highest positive correlation was yielded by the `NumGap` feature of B_6

[1]All of the features are computed for each of the behaviours but Episodes, which is behaviour agnostic.

Table 5: Behaviours and corresponding features correlated with achievement, where correlation coefficients $| > 0.45|$ and $p < 0.001$.

Id	Interactive behaviour	Correlated feature	Spearman's ρ	Kendall's τ
B_1	Explore the week 1 page main area, left sidebar, and forum main area	NumGap	0.55	0.54
		Epi_not0	0.47	0.45
B_2	Explore the week 1 page main area then leave page	S10mean	-0.48	-0.47
		T10mean	-0.50	-0.49
B_3	Explore the week 1 main area and forum main area	NumGap	0.48	0.46
B_4	Explore the homepage main area and interact with videos	NumGap	0.49	0.46
B_5	Explore the forum main area and left sidebar	NumGap	0.57	0.55
		Sum	0.51	0.49
		Pl10sum	-0.49	-0.46
		Pl5sum	-0.49	-0.46
B_6	Explore the forum main area	NumGap	0.61	0.57
		Epi_not0	0.58	0.53
		S3sum	-0.72	-0.71
		S5sum	-0.64	-0.63

($\rho = 0.61$, $p < 0.001$), which indicates a relationship between obtaining a badge and the number of continuous sessions where the forum's main area was not explored. On the other hand, the highest negative correlation was `S3sum` ($\rho = -0.72$, $p < 0.001$) for the same behaviour, suggesting a relationship between achievement and a decrease in the exploration of the forum over time.

Table 6: Comparison between models including all the features (`AllFeatures`) and models including engagement features only (`EngFeatures`)

	Cox & Snell R^2			**Nagelkerke R^2**		
	`AllFeatures`	`EngFeatures`	Diff	`AllFeatures`	`EngFeatures`	Diff
Best	0.43	0.38	+0.05	0.82	0.73	+0.09
Average	0.42	0.35	+0.07	0.81	0.67	+0.14
Median	0.42	0.36	+0.06	0.81	0.69	+0.12

For the six interactive behaviours that yielded coefficients above 0.45 with badge status in Table 5, we included in a regression analysis the top two strongly correlated positive/negative features of each behaviour as long as they were not variations of the same feature (e.g. `S3sum` and `S5sum` for B_6) to prevent overfitting. These were the dependent variables of the regression analysis, while badge status was the dependent variable. As dependent variable is binary, we used logistic regression analysis. The logistic models with 1–6 predictors were computed following the stepwise backward and forward techniques, where `Model 1` has the highest explainability (81.8%, Nagelkerke R^2) about whether a student received a badge:

$$-9.44 + 1.52B_1 - 1.08B_2 - 1.15B_3 - 0.04B_4 - 1.06B_5 + 2.40B_6 \quad (1)$$

The goodness-of-fit test (Hosmer–Lemeshow test) yielded a value of 0.246 and was insignificant ($p = 0.993$), suggesting that the model fits the data well. Two descriptive measures of goodness-of-fit presented are Cox and Snell $R^2 = 0.43$ as well as Nagelkerke. These indices are variations of the R^2 concept defined for Ordinary Least Squares regression (OLS) model.

According to `Model 1`, the log of the odds of receiving a badge is negatively associated with the exploration and withdrawal of the week 1 page main area ($B_2, p = 0.01$), the exploration of the week 1 main area followed by the exploration of the forum ($B_3, p = 0.37$), the exploration of the homepage followed by the interaction with videos ($B_4, p = 0.96$), and the exploration of the forum main area and left sidebar ($B_5, p = 0.36$). Those behaviours that were positively associated with receiving a badge include the exploration of the week 1 page main area, left sidebar, and the forum main area ($B_1, p = 0.30$), and the exploration of the forum main area ($B_6, p = 0.01$). The statistical significance of individual regression coefficients were tested using the Wald chi-square statistic. According to the results, only B_2's `S10mean` ($p = 0.01$) and B_6's `NumGap` ($p = 0.01$) were significant predictors of badge status.

In other words, the more students revisited (`S10mean`) week 1 materials to then leave the learning platform without interacting with any other section of the MOOC (B_2), the less likely they were to receive a badge. This could be because this behaviour models those who got stuck in week 1, try to catch up several times but make no further advances in the course. We confirm that the number of consecutive sessions in which the forum was not explored (`NumGap` on B_6) is a strong predictor of badge status. The odds of students exhibiting B_2 (`S10mean`) not receiving a badge were 3, whereas the odds for those exhibiting B_6 (`NumGap`) receiving a badge were 11.

Some of the features included in our previous regression analysis are closely related to engagement, including the number of inactive sessions. We know from previous work that dwell time is related to engagement [88, 53, 41, 78]. To investigate the added value of including non-engagement metrics in the models and address **RQ3**, we computed a baseline model with those features in Table 3 that are known indicators of engagement. Similar to the previous analysis, the variable that contributed most was the number of consecutive sessions (`NumGap`) in which the forum was not explored (B_6). The model with the highest explanatory power (Nagelkerke $R^2 = 72.8\%$) for whether a given student receives a badge is:

$$-6.52 + 0.36B_1 - 0.001B_3 - 0.23B_4 - 0.62B_5 + 1.85B_6 \tag{2}$$

where only B_6 is a significant predictor of achievement ($p = 0.001$). We computed R^2 values of the top-20 performing models for both analyses: i.e. the models including all features (`AllFeatures`) and the models including only engagement features (`EngFeatures`) —see Table 6. Our results indicate that including non-engagement behaviours in the models adds a 9% of explainability in the best case and a 14% on average for the Nagelkerke R^2.

Table 7: 10-fold cross validation of the original dataset.

Classifiers	**Precision**	**Recall**	**F1 Score**	**Accuracy**	**Accuracy Change**
`K Nearest Neighbour`	0.65	0.60	0.62	0.94	+0.06
`Naive Bayes`	0.72	0.80	0.73	0.92	+0.04
`Decision Tree`	0.72	0.75	0.69	0.93	+0.05
`Random Forest`	0.85	0.75	0.78	0.94	+0.06
`Logistic Regression`	0.77	0.75	0.73	0.94	+0.06
`Support Vector Machine`	0.90	0.55	0.67	0.94	+0.06
`AdaBoost`	0.48	0.50	0.48	0.91	+0.03
`Multi-layer Perception`	0.57	0.60	0.56	0.93	+0.05
`Bagging`	0.90	0.60	0.70	0.95	+0.07
`Gradient Tree Boosting`	0.75	0.60	0.65	0.94	+0.06

Table 8: 10-fold cross validation with oversampling.

Classifiers	Precision	Recall	F1 Score	Accuracy	Accuracy Change
K Nearest Neighbour	0.58	0.85	0.66	0.89	+0.01
Naive Bayes	0.57	0.95	0.69	0.89	+0.01
Decision Tree	0.78	0.70	0.69	0.93	+0.05
Random Forest	0.75	0.75	0.71	0.92	+0.04
Logistic Regression	0.73	0.85	0.73	0.91	+0.03
Support Vector Machines	0.63	0.85	0.71	0.91	+0.03
AdaBoost	0.72	0.80	0.73	0.94	+0.06
Multi-layer Perception	0.66	0.85	0.71	0.93	+0.05
Bagging	0.57	0.90	0.65	0.89	+0.01
Gradient Tree Boosting	0.62	0.65	0.59	0.92	+0.04

We then conducted 10-fold cross validation using a variety of classifiers. As the distribution of students with and without a badge is imbalanced, only 12% (17 out of 140) of the participants received a badge with a 1:7.2 ratio, to avoid losing information in majority examples to undersampling [42] and overly-optimistic estimates [76], we performed oversampling during the cross validation procedure. Synthetic Minority Oversampling Technique (SMOTE) coupled with Tomek Links was performed as the combination of algorithms was suggested to prevent overfitting [76]. Stratified k-fold cross validation was used to ensure that the proportion of positive to negative examples is kept in the folds. The results for the 10-fold cross validation on the original dataset and with oversampling are shown in Tables 7 and 8, where the accuracy change (difference between the accuracy of the model and the default distribution percentage which is the students without a badge out of all the students), precision, recall, and f1 score are tabulated. We focus on the precision and accuracy as the importance of identifying students with tendencies to fail is greater than identifying students who are likely to succeed. For the original dataset (without oversampling), the accuracy of the models identifying failing students is 91-95%, 3-7% higher than using random selection (i.e. the percentage of students without a badge, 88%). The Bagging classifier yields the largest increase in accuracy across the two sets of results (7%), with precision being 0.90.

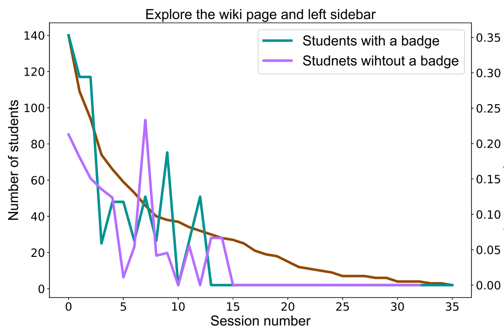

(a) Cluster 1

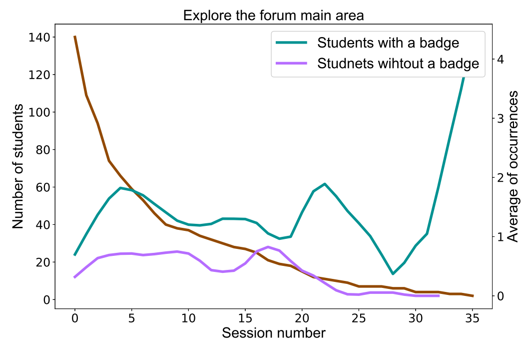

(b) Cluster 2

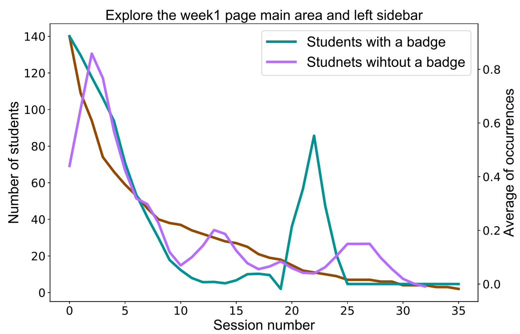

(c) Cluster 3

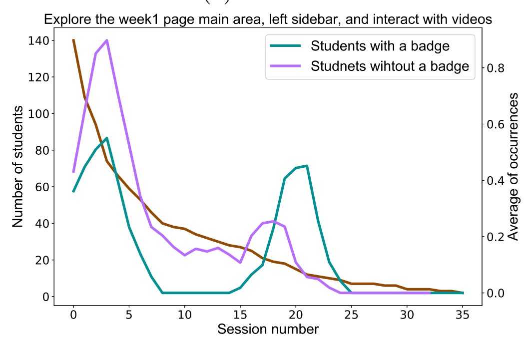

(d) Cluster 4

Figure 1: Graphical representation of the number of the active users for each session (brown line), the average of occurrences for each session among students who received a badge (blue line), the average of occurrences for each session among students who did not receive a badge (purple line).

To investigate the occurrences of the behaviours over time, we analysed the occurrences of the 23 behaviours in each session for the students that eventually received a badge and students who did not. The number of the active students for each session was plotted alongside the average of occurrences for each session among the students who received a badge and the students who did not (data had been smoothed using LOESS Curve Fitting). Based on our observations of the visualisations as shown in Figure 1, we classified these behaviours into four clusters:

1. Behaviours of both groups evolve similarly over time (Figure 1a).
2. Behaviours were exhibited more by students with the badge (Figure 1b). Among the six behaviours included in our models, B_1, B_3, B_5, B_6 are in this cluster.

3. Behaviours of both groups evolve similarly for most sessions, with significant differences only at specific sessions (Figure 1c). Behaviours B_2 and B_4 are classified in this cluster.
4. Behaviours were exhibited more by students without the badge for most sessions, with significantly higher peak for badge holders in specific sessions (Figure 1d).

6. Discussion

We revisit the research questions formulated at the outset:

RQ1. *Can we identify micro-interactions that are relevant for online learning?*
We utilised sequential pattern mining and thematic analysis to find implicit regularities that are associated with learning including achievement and unit progress. Our methodology allowed us to investigate micro-interactions in different abstraction levels by grouping fine-grained interaction patterns into higher level behaviours. Our findings suggest that micro-interactions can potentially be used to predict achievement.

RQ2. *Which micro-interactions are associated with learning progress and achievement?*
We confirmed existing findings on an online learning platform and discovered that interactive behaviours involving exploration of forum, and leaving the online platform after the exploration of the week 1 materials are significant predictors of student achievement. Further to this, our findings provide more clarity on these behaviours: first, the less students abandon week 1 materials, the more likely they are to receive a badge; second, the number of consecutive sessions where the forum's main area was not explored was strongly related to obtaining a badge. While the first behaviour confirms that withdrawal is directly related to the (lack of) achievement, it is unexpected that the less the forum is explored the more likely students are to achieve a badge [7, 19, 71]. This may be due to the forum being perceived as a collaborative tool within the context of a cMOOC, which was used as a resource for students who were struggling. Hence, not required by advanced students.

The interactive behaviours include regularities such as exploration of the forum and interactions with videos. Forum participation and video interactions have been thoroughly investigated in previous research [75, 82, 10, 49], perhaps as they are the most representative of a classic classroom structure. The other frequent elements, for example exploration of homepage, week 1 materials, and leaving week 1 materials are within expectation: exploring the week 1 materials and leaving may be interpreted as distraction from the learning material and abandonment of the course, which resulted in low achievement. This is in line with the visualisations in Figure 1d illustrating the decrease of interactions over time and previous research indicating that participation tends to drop rapidly within the first week [28].

We found that a decrease in exploring the forum is a strong predictor of achievement while the visualisations show that successful students were more active in the forum. The findings may seem contradictory as they indicate that the levels of participation in forums decreased over time for students who received a badge while still remained higher compared to other students throughout the sessions. More forum activities suggests that the student is more likely to receive a badge, which is in line with related studies showing that forum participation positively affects the learning outcome [7, 19, 71]. Our findings can be explained in that, in sessions when students were less active in the forum, they were conducting other learning activities in the platform. It may also be the case that as students use the forum to make progress and learn, they depend on the forum less and become more autonomous learners.

We did not find strong relationships with the number of units completed. This may be because obtaining the badge involves a set of diverse activities and it is more informative than student progress. The models can be used to develop, for instance, browser extensions to capture interface interactions in real time and predict student learning trajectory to provide feedback and guidance to students and course leaders.

RQ3. *What is the added value of using micro-interactions?*
Features of engagement (e.g. the number of inactive sessions) are the strongest predictor of achievement according to our analysis. The explainability of models built with engagement features is as high as 72.8%. By comparing the models from two regression analysis containing different features, we discovered a 10% contribution to the explainability of student achievement from features of behaviours other than engagement, and perhaps learnability. When analysing the performance of the models, results indicate a 3-7% increase in accuracy in identifying students who will not receive a badge versus random selection (i.e. the percentage of students who did not re-

ceive a badge) suggesting that micro-interactions implicitly capture learning processes. This can be explained in that the relationship between cognitive engagement and learning outcomes is mediated by user activity [32]. Consequently, it can be argued that the user activity in online learning platforms contains behavioural markers that are indicators of cognitive processes and learning outcome.

6.1. The Role of User Interface Learnability

In general, beyond the decrease in the interactions of successful students with the forum, we discovered that students exhibited decreasing trends of participation throughout the online sessions. The decrease can be expected as engagement typically decreases over time on online learning platforms. The correlation results showed that most of the trends are negatively correlated with student achievement, which means that the less a student exhibits these behaviours, the more likely they are to receive a badge. As the trend features were calculated using windows of 3/5/10 sessions, the trends can be considered as fluctuations, which is confirmed by the visualisations of occurrences, as shown in Figure 1a. More fluctuations may mean that students exhibited different interactive behaviours at the same time. As students engage with the platform they become more familiar with the platform, and consequently they become more proficient at navigating and using it. This finding may further support our assumption that interface learnability has an effect on the learning outcome as measured by student achievement. Usability has been investigated in online learning settings to evaluate different platforms and to provide insights on the design of platforms and the learning resources [87, 80]. Although it has been suggested that there is no widely-agreed definition for learnability, it can be described as 'the ability to improve performance over entire usage history' based on a taxonomy provided in [39]. Learnability as an important and perhaps fundamental component of usability [2, 60], has not been thoroughly discussed in online learning platforms.

The plausible association between *becoming a more efficient user* and student achievement may suggest that the easier a learning platform is to use, the better students learn. It may also suggest that students who are more familiar with the learning platform or skilled navigators are better at learning. This opens up new research avenues into analysing the learnability of platforms and student learning outcome.

6.2. Implications

Methodological implications. Previous research has identified the challenges of using sequential pattern mining algorithms in educational studies [68]: 1) the generation of excessive patterns with limited relevancy and value, and 2) the involvement of domain experts for filtering and labeling purposes. For the first challenge, we proposed a benchmarking process to select the most suitable algorithm, dataset and algorithm parametrisation so that we maximise how informative and representative the results are. The second challenge is relatively common in data mining and analysis scenarios. Through a number of established methodological steps, thematic analysis facilitates the work of domain experts.

Theoretical implications. We found that micro-interactions can be behavioural markers of online learning. Whether these micro-interactions are universal is unknown but they are probably not generalisable across online platforms. However, they provide further explainability in terms of interpretability and predictability of what students do.

Implications for practice. Despite the current demands on the use of rich interaction data [56], existing data analysis procedures in online learning do not contemplate low-level events. This may be due to the lack of infrastructure to do so although, actually, there is an availability of these tools [8] which facilitate low-level data collection. However, computing patterns and building student models require programming and data science expertise. While they are not widespread and their coverage is limited, tools such as WevQuery-PM [9] can lighten this burden.

6.3. Limitations

As each online learning platform is designed for different purposes and audiences, there may be limitations to the extent of the generalisability of our conclusions, which a common issue in online learning [4]. It is well known that the models defined for MOOCs are typically valid within the scope of the MOOCs under scrutiny: early career researchers on a cMOOC. There are no universal models as students, learning contents, and course designs differ [58]. Hence, researchers emphasise on methodological approaches to build student models [52]. We contribute to such body of knowledge with this novel approach.

Also, compared to other online learning platforms, the sample size is relatively small due to the fact that the platform was targeting an specialised student group, early career researchers. We counterbalance this limitation with the large amount of data points collected, which give us an in-depth understanding of student behaviour.

While students interactions with the platform are used to predict whether they will receive a badge or not, achieving a badge is also rewarded partly for these same interactions. It could be argued, however, that low-level interactions used in analysis contain many more varieties of higher level interactions (such as viewing the video content and wiki page) than those used to give out the badge. Further to this, the criteria for receiving a badge includes activities and assignments which are not captured in these low-level interactions.

7. Conclusion

We propose a three-step methodology to determine if there are implicit micro-interactions that are associated with indicators of learning. To suggest its feasibility, we apply the methodology in an online learning platform. From the generated dataset, we extracted frequent student interaction patterns using sequential pattern mining algorithms. To handle the sheer numbers of patterns we applied thematic analysis to group the patterns into themes that represent interactive behaviours. Then we sought the presence of these representative behaviours in the original dataset to build student models that represent such behaviours. Finally, we conducted statistical analysis between features derived from the models and indicators of learning including achievement and unit progress.

We identified features from six interactive behaviours that strongly correlated with achievement that explain 72% of the student achievement variation, which was 10% higher when non engagement features were included. The models with the interactive behaviours were 3-7% more accurate at identifying students who will not receive a badge than random selection. It would have not been possible to obtain the increases in explainability and accuracy without our data-driven approach. This increase involves deeper insights gained from the relationship between student achievement and learnability of the student interface. In summary, our data-driven approach provides self-regulated learning platforms with a fresh perspective on measuring indicators of learning, and has revealed connections between platform/student learnability and learning outcome.

Acknowledgement

This work was supported by the EU's Horizon 2020 programme under grant agreement H2020-693092 MOVING (http://moving-project.eu).

We thank Emmanuel Wolde-Michael for his help and support with thematic analysis.

References

[1] Abramovich, S., Schunn, C., Higashi, R., 2013. Are badges useful in education?: It depends upon the type of badge and expertise of learner. Educational Technology Research and Development 61. doi:10.1007/s11423-013-9289-2.

[2] Abran, A., Khelifi, A., Suryn, W., Seffah, A., 2003. Usability meanings and interpretations in ISO standards, in: Software Quality Journal, Kluwer Academic Publishers. pp. 325–338. doi:10.1023/A:1025869312943.

[3] Ali, F., George, A., 2015. Impact of a formative e-assessment on learning outcomes: A pilot study on a social and behavioural sciences course, college of health sciences, university of bahrain, in: 2015 Fifth International Conference on e-Learning (econf), pp. 408–412. doi:10.1109/ECONF.2015.38.

[4] Alonso-Mencía, M.E., Alario-Hoyos, C., Maldonado-Mahauad, J., Estévez-Ayres, I., Pérez-Sanagustín, M., Kloos, C.D., 2020. Self-regulated learning in moocs: lessons learned from a literature review. Educational Review 72, 319–345. doi:10.1080/00131911.2019.1566208.

[5] Aloysius, G., Binu, D., 2013. An approach to products placement in supermarkets using PrefixSpan algorithm. Journal of King Saud University - Computer and Information Sciences 25, 77–87. doi:10.1016/j.jksuci.2012.07.001.

[6] Alturkistani, A., Majeed, A., Car, J., Brindley, D., Wells, G., Meinert, E., 2019. Data collection approaches to enable evaluation of a massive open online course about data science for continuing education in health care: Case study. Journal of Medical Internet Research 21. doi:10.2196/10982.

[7] Anderson, A., Huttenlocher, D., Kleinberg, J., Leskovec, J., 2014. Engaging with massive online courses, in: Proceedings of the 23rd international conference on World wide web - WWW '14, ACM Press, New York, New York, USA. pp. 687–698. doi:10.1145/2566486.2568042.

[8] Apaolaza, A., Vigo, M., 2017. Wevquery: Testing hypotheses about web interaction patterns. Proc. ACM Hum.-Comput. Interact. 1. URL: https://doi.org/10.1145/3095806, doi:10.1145/3095806.

[9] Apaolaza, A., Vigo, M., 2019. Assisted pattern mining for discovering interactive behaviours on the web. International Journal of Human-Computer Studies 130, 196 – 208. URL: https://doi.org/10.1016/j.ijhcs.2019.06.012.

[10] Atapattu, T., Falkner, K., 2017. Discourse analysis to improve the effective engagement of mooc videos, in: Proceedings of the Seventh International Learning Analytics amp; Knowledge Conference, Association for Computing Machinery, New York, NY, USA. p. 580–581. doi:10.1145/3027385.3029470.

[11] Baek, J., Shore, J., 2016. Promoting Student Engagement in MOOCs, in: Proceedings of the Third (2016) ACM Conference on Learning @ Scale - L@S '16, ACM Press, New York, New York, USA. pp. 293–296. doi:10.1145/2876034.2893437.

[12] de Barba, P., Kennedy, G., Ainley, M., 2016. The role of students' motivation and participation in predicting performance in a mooc. Journal of Computer Assisted Learning 32, 218–231. doi:https://doi.org/10.1111/jcal.12130.

[13] Braun, V., Clarke, V., 2006. Using thematic analysis in psychology. Qualitative Research in Psychology 3, 77–101. doi:10.1191/1478088706qp063oa.

[14] Breslav, S., Khan, A., Hornbæk, K., 2014. Mimic: Visual analytics of online micro-interactions, in: Proceedings of the Workshop on Advanced Visual Interfaces AVI, Association for Computing Machinery. pp. 245–252. doi:10.1145/2598153.2598168.

[15] Breslow, L., Pritchard, D.E., DeBoer, J., Stump, G.S., Ho, A.D., Seaton, D.T., 2013. Studying Learning in the Worldwide Classroom Research into edX's First MOOC. Research & Practice in Assessment 8, 13–25. URL: https://eric.ed.gov/?id=EJ1062850.

[16] Brooks, C., Thompson, C., Teasley, S., 2015. A time series interaction analysis method for building predictive models of learners using log data, in: Proceedings of the Fifth International Conference on Learning Analytics And Knowledge, Association for Computing Machinery, New York, NY, USA. p. 126–135. doi:10.1145/2723576.2723581.

[17] Coleman, C.A., Seaton, D.T., Chuang, I., 2015. Probabilistic use cases: Discovering behavioral patterns for predicting certification, in: L@S 2015 - 2nd ACM Conference on Learning at Scale, Association for Computing Machinery, Inc. pp. 141–148. doi:10.1145/2724660.2724662.

[18] Cormier, D., Siemens, G., 2010. Through the open door: Open courses as research, learning, and engagement. EDUCAUSE Review 45, 30–39.

[19] Crues, R.W., Bosch, N., Perry, M., Angrave, L., Shaik, N., Bhat, S., 2018. Refocusing the lens on engagement in MOOCs, in: Proceedings of the Fifth Annual ACM Conference on Learning at Scale - L@S '18, ACM Press, New York, New York, USA. pp. 1–10. doi:10.1145/3231644.3231658.

[20] Daradoumis, T., Martínez-Monés, A., Xhafa, F., 2006. A layered framework for evaluating on-line collaborative learning interactions. International Journal of Human-Computer Studies 64, 622 – 635. doi:https://doi.org/10.1016/j.ijhcs.2006.02.001.

[21] DeBoer, J., Breslow, L., 2014. Tracking progress: Predictors of students' weekly achievement during a circuits and electronics mooc, in: Proceedings of the First ACM Conference on Learning @ Scale Conference, Association for Computing Machinery, New York, NY, USA. p. 169–170. doi:10.1145/2556325.2567863.

[22] Deeva, G., De Smedt, J., De Koninck, P., Weerdt, J., 2018. Dropout prediction in moocs: A comparison between process and sequence mining, pp. 243–255. doi:10.1007/978-3-319-74030-0_18.

[23] Dev, H., Liu, Z., 2017. Identifying frequent user tasks from application logs, in: Proceedings of the 22Nd International Conference on Intelligent User Interfaces, ACM, New York, NY, USA. pp. 263–273. doi:10.1145/3025171.3025184.

[24] Ding, L., Er, E., Orey, M., 2018. An exploratory study of student engagement in gamified online discussions. Computers Education 120, 213 – 226. doi:https://doi.org/10.1016/j.compedu.2018.02.007.

[25] Doko, E., Abazi Bexheti, L., Hamiti, M., Prevalla Etemi, B., 2018. Sequential Pattern Mining Model to Identify the Most Important or Difficult Learning Topics via Mobile Technologies. International Journal of Interactive Mobile Technologies (iJIM) 12, 109. doi:10.3991/ijim.v12i4.9223.

[26] Dunaev, M., Zaytsev, K., 2017. Using sequential pattern mining for building efficient massive open online courses. International Journal of Pure and Applied Mathematics 117, 131–135.

[27] Eickhoff, C., Teevan, J., White, R., Dumais, S., 2014. Lessons from the journey, in: Proceedings of the 7th ACM international conference on Web search and data mining - WSDM '14, ACM Press, New York, New York, USA. pp. 223–232. doi:10.1145/2556195.2556217.

[28] Evans, B.J., Baker, R.B., Dee, T.S., 2016. Persistence patterns in massive open online courses (moocs). The Journal of Higher Education 87, 206–242. doi:10.1080/00221546.2016.11777400.

[29] Falakmasir, M.H., Gonzalez-Brenes, J.P., Gordon, G.J., DiCerbo, K.E., 2016. A Data-Driven Approach for Inferring Student Proficiency from Game Activity Logs, in: Proceedings of the Third (2016) ACM Conference on Learning @ Scale - L@S '16, ACM Press, New York, New York, USA. pp. 341–349. doi:10.1145/2876034.2876038.

[30] Farag, W., 2012. Comparing achievement of intended learning outcomes in online programming classes with blended offerings, in: Proceedings

of the 13th Annual Conference on Information Technology Education, Association for Computing Machinery, New York, NY, USA. p. 25–30. doi:10.1145/2380552.2380561.

[31] Ferguson, R., 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, Association for Computing Machinery, New York, NY, USA. p. 51–58. doi:10.1145/2723576.2723606.

[32] Fincham, E., Whitelock-Wainwright, A., Kovanović, V., Joksimović, S., van Staalduinen, J.P., Gašević, D., 2019. Counting clicks is not enough: Validating a theorized model of engagement in learning analytics, in: Proceedings of the 9th International Conference on Learning Analytics amp; Knowledge, Association for Computing Machinery, New York, NY, USA. p. 501–510. doi:10.1145/3303772.3303775.

[33] Fournier Viger, P., Lin, C.W., Rage, U., Koh, Y.S., Thomas, R., 2017. A survey of sequential pattern mining. Data Science and Pattern Recognition 1, 54–77.

[34] Fournier Viger, P., Wu, C.W., Gomariz, A., Tseng, V., 2014. Vmsp: Efficient vertical mining of maximal sequential patterns. doi:10.1007/978-3-319-06483-3_8.

[35] Fredricks, J.A., Blumenfeld, P.C., Paris, A.H., 2004. School Engagement: Potential of the Concept, State of the Evidence. Review of Educational Research 74, 59–109. doi:10.3102/00346543074001059.

[36] Geigle, C., Zhai, C., 2017. Modeling MOOC Student Behavior With Two-Layer Hidden Markov Models, in: Proceedings of the Fourth (2017) ACM Conference on Learning @ Scale - L@S '17, ACM Press, New York, New York, USA. pp. 205–208. doi:10.1145/3051457.3053986.

[37] Gomez Zermeño, M., Aleman de la Garza, L., 2016. Research analysis on mooc course dropout and retention rates. Turkish Online Journal of Distance Education 17. doi:10.17718/tojde.23429.

[38] Gong, L., Liu, Y., Zhao, W., 2018. Using learning analytics to promote student engagement and achievement in blended learning: An

empirical study, in: Proceedings of the 2nd International Conference on E-Education, E-Business and E-Technology, Association for Computing Machinery, New York, NY, USA. p. 19–24. doi:10.1145/3241748.3241760.

[39] Grossman, T., Fitzmaurice, G., Attar, R., 2009. A survey of software learnability: Metrics, methodologies and guidelines, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Association for Computing Machinery, New York, NY, USA. p. 649–658. doi:10.1145/1518701.1518803.

[40] Guerra, J., Sahebi, S., Lin, Y.R., Brusilovsky, P., 2014. The problem solving genome: Analyzing sequential patterns of student work with parameterized exercises .

[41] Guo, P.J., Kim, J., Rubin, R., 2014. How video production affects student engagement: An empirical study of mooc videos, in: Proceedings of the First ACM Conference on Learning @ Scale Conference, Association for Computing Machinery, New York, NY, USA. p. 41–50. doi:10.1145/2556325.2566239.

[42] Haixiang, G., Li, Y., Shang, J., Mingyun, G., Yuanyue, H., Gong, B., 2016. Learning from class-imbalanced data: Review of methods and applications. Expert Systems with Applications 73. doi:10.1016/j.eswa.2016.12.035.

[43] Hamari, J., 2017. Do badges increase user activity? a field experiment on the effects of gamification. Computers in Human Behavior 71, 469 – 478. doi:https://doi.org/10.1016/j.chb.2015.03.036.

[44] Hilbert, D.M., Redmiles, D.F., 2000. Extracting usability information from user interface events. ACM Computing Surveys 32, 384–421. doi:10.1145/371578.371593.

[45] Hosseini, R., Brusilovsky, P., Yudelson, M., Hellas, A., 2017. Stereotype modeling for problem-solving performance predictions in moocs and traditional courses, in: Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization, Association for Computing Machinery, New York, NY, USA. p. 76–84. doi:10.1145/3079628.3079672.

[46] Jansen, B.J., Booth, D., Smith, B., 2009. Using the taxonomy of cognitive learning to model online searching. Information Processing & Management 45, 643–663. doi:10.1016/j.ipm.2009.05.004.

[47] Jones, R., Klinkner, K.L., 2008. Beyond the session timeout: Automatic hierarchical segmentation of search topics in query logs, in: Proceedings of the 17th ACM Conference on Information and Knowledge Management, Association for Computing Machinery, New York, NY, USA. p. 699–708. doi:10.1145/1458082.1458176.

[48] Käser, T., Hallinen, N.R., Schwartz, D.L., 2017. Modeling exploration strategies to predict student performance within a learning environment and beyond, in: Proceedings of the Seventh International Learning Analytics & Knowledge Conference on - LAK '17, ACM Press, New York, New York, USA. pp. 31–40. doi:10.1145/3027385.3027422.

[49] Kim, J., Li, S.W., Cai, C.J., Gajos, K.Z., Miller, R.C., 2014. Leveraging video interaction data and content analysis to improve video learning, in: Proceedings of the CHI 2014 Learning Innovation at Scale workshop.

[50] Kizilcec, R.F., Piech, C., 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, Association for Computing Machinery, New York, NY, USA. p. 170–179. doi:10.1145/2460296.2460330.

[51] Kuh, G.D., Cruce, T.M., Shoup, R., Kinzie, J., Gonyea, R.M., 2008. Unmasking the effects of student engagement on first-year college grades and persistence. The Journal of Higher Education 79, 540–563. doi:10.1080/00221546.2008.11772116.

[52] Lehmann, J., Lalmas, M., Yom-Tov, E., Dupret, G., 2012. Model of user engagement. doi:10.1007/978-3-642-31454-4_14.

[53] Lu, H., Zhang, M., Ma, W., Shao, Y., Liu, Y., Ma, S., 2019. Quality effects on user preferences and behaviors in mobile news streaming, in: The World Wide Web Conference, ACM, New York, NY, USA. pp. 1187–1197. doi:10.1145/3308558.3313751.

[54] Lushnikova, N., Chintakayala, P., Rodante, A., 2012. Massive Open Online Courses from Ivy League Universities: Benefits and Challenges for Students and Educators.

[55] Maldonado, J., Pérez-Sanagustín, M., Moreno-Marcos, P., Alario-Hoyos, C., Merino, P., Delgado-Kloos, C., 2018. Predicting Learners' Success in a Self-paced MOOC Through Sequence Patterns of Self-regulated Learning: 13th European Conference on Technology Enhanced Learning, EC-TEL 2018, Leeds, UK, September 3-5, 2018, Proceedings. pp. 355–369. doi:10.1007/978-3-319-98572-5_27.

[56] Maldonado-Mahauad, J., Pérez-Sanagustín, M., Kizilcec, R.F., Morales, N., Munoz-Gama, J., 2018. Mining theory-based patterns from big data: Identifying self-regulated learning strategies in massive open online courses. Computers in Human Behavior 80, 179 – 196. doi:https://doi.org/10.1016/j.chb.2017.11.011.

[57] Meinert, E., Alturkistani, A., Brindley, D., Carter, A., Wells, G., Car, J., 2018. Protocol for a mixed-methods evaluation of a massive open online course on real world evidence. BMJ Open 8. doi:10.1136/bmjopen-2018-025188.

[58] Motz, B., Quick, J., Schroeder, N., Zook, J., Gunkel, M., 2019. The validity and utility of activity logs as a measure of student engagement, in: ACM International Conference Proceeding Series, Association for Computing Machinery. pp. 300–309. doi:10.1145/3303772.3303789.

[59] Muldner, K., Burleson, W., Sande, B., Vanlehn, K., 2011. An analysis of students' gaming behaviors in an intelligent tutoring system: Predictors and impacts. User Model. User-Adapt. Interact. 21, 99–135. doi:10.1007/s11257-010-9086-0.

[60] Nielsen, J., 1994. Usability Engineering. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

[61] Oliphant, T.E., 2006. A guide to NumPy. volume 1. Trelgol Publishing USA.

[62] Onah, D., Sinclair, J., Boyatt, R., 2014. Dropout rates of massive open online courses: Behavioural patterns doi:10.13140/RG.2.1.2402.0009.

[63] Ossiannilsson, E., Altinay, F., Altinay, Z., 2015. Analysis of MOOCs practices from the perspective of learner experiences and quality culture. Educational Media International 52, 272–283. doi:10.1080/09523987.2015.1125985.

[64] Papathoma, T., Blake, C., Clow, D., Scanlon, E., 2015. Investigating Learners' Views of Assessment Types in Massive Open Online Courses (MOOCs). pp. 617–621. doi:10.1007/978-3-319-24258-3_72.

[65] Pardos, Z.A., Tang, S., Davis, D., Le, C.V., 2017. Enabling Real-Time Adaptivity in MOOCs with a Personalized Next-Step Recommendation Framework, in: Proceedings of the Fourth (2017) ACM Conference on Learning @ Scale - L@S '17, ACM Press, New York, New York, USA. pp. 23–32. doi:10.1145/3051457.3051471.

[66] Park, J., Denaro, K., Rodriguez, F., Smyth, P., Warschauer, M., 2017. Detecting changes in student behavior from clickstream data, in: Proceedings of the Seventh International Learning Analytics & Knowledge Conference on - LAK '17, ACM Press, New York, New York, USA. pp. 21–30. doi:10.1145/3027385.3027430.

[67] Perera, D., Kay, J., Koprinska, I., Yacef, K., Zaiane, O., 2008. Clustering and sequential pattern mining of online collaborative learning data. IEEE Transactions On Knowledge and Data Engineering 21, 759–772.

[68] Poon, L., Kong, S.c., Wong, M., Yau, T., 2017. Mining sequential patterns of students' access on learning management system, pp. 191–198. doi:10.1007/978-3-319-61845-6_20.

[69] Prince, M., 2004. Does Active Learning Work? A Review of the Research. Journal of Engineering Education 93, 223–231. doi:10.1002/j.2168-9830.2004.tb00809.x.

[70] Pérez-Sanagustín, M., Sharma, K., Pérez, R., Maldonado, J., Broisin, J., 2019. Analyzing Learners' Behavior Beyond the MOOC: An Exploratory Study. pp. 40–54. doi:10.1007/978-3-030-29736-7_4.

[71] Ramesh, A., Goldwasser, D., Huang, B., Daume, H., Getoor, L., 2014. Uncovering hidden engagement patterns for predicting learner performance in MOOCs, in: Proceedings of the first ACM conference on

Learning @ scale conference - L@S '14, ACM Press, New York, New York, USA. pp. 157–158. doi:10.1145/2556325.2567857.

[72] Robinson, C., Yeomans, M., Reich, J., Hulleman, C., Gehlbach, H., 2016. Forecasting student achievement in moocs with natural language processing, pp. 383–387. doi:10.1145/2883851.2883932.

[73] Rodriguez, C.O., 2012. Moocs and the ai-stanford like courses: Two successful and distinct course formats for massive open online courses. The European Journal of Open, Distance and E-Learning 15.

[74] Rohloff, T., Sauer, D., Meinel, C., 2020. Students' achievement of personalized learning objectives in moocs, in: Proceedings of the Seventh ACM Conference on Learning @ Scale, Association for Computing Machinery, New York, NY, USA. p. 147–156. doi:10.1145/3386527.3405918.

[75] Romero, C., López, M.I., Luna, J.M., Ventura, S., 2013. Predicting students' final performance from participation in on-line discussion forums. Computers & Education 68, 458–472. doi:10.1016/j.compedu.2013.06.009.

[76] Santos, M., Soares, J., Henriques Abreu, P., Araujo, H., Santos, J., 2018. Cross-validation for imbalanced datasets: Avoiding overoptimistic and overfitting approaches. IEEE Computational Intelligence Magazine 13, 59–76. doi:10.1109/MCI.2018.2866730.

[77] Shen, H., Liang, L., Law, N., Hemberg, E., O'Reilly, U.M., 2020. Understanding learner behavior through learning design informed learning analytics, in: Proceedings of the Seventh ACM Conference on Learning @ Scale, Association for Computing Machinery, New York, NY, USA. p. 135–145. doi:10.1145/3386527.3405919.

[78] Singh, V., Padmanabhan, B., de Vreede, T., de Vreede, G.J., Andel, S., Spector, P.E., Benfield, S., Aslami, A., 2018. A content engagement score for online learning platforms, in: Proceedings of the Fifth Annual ACM Conference on Learning at Scale - L@S '18, ACM Press, New York, New York, USA. pp. 1–4. doi:10.1145/3231644.3231683.

[79] Sun, J.C.Y., Rueda, R., 2012. Situational interest, computer self-efficacy and self-regulation: Their impact on student engagement in distance

education. British Journal of Educational Technology 43, 191–204. doi:10.1111/j.1467-8535.2010.01157.x.

[80] Tsironis, A., Katsanos, C., Xenos, M., 2016. Comparative usability evaluation of three popular mooc platforms, in: 2016 IEEE Global Engineering Education Conference (EDUCON), pp. 608–612. doi:10.1109/EDUCON.2016.7474613.

[81] Wang, W., Weng, J.F., Su, J.M., Tseng, S.S., 2004. Learning portfolio analysis and mining in scorm compliant environment, pp. T2C – 17. doi:10.1109/FIE.2004.1408490.

[82] Wang, X., Yang, D., Wen, M., Koedinger, K., Rosé, C.P., 2015. Investigating how student's cognitive behavior in mooc discussion forums affect learning gains. International Educational Data Mining Society .

[83] Wham, D., 2017. Forecasting student outcomes at university-wide scale using machine learning, in: Proceedings of the Seventh International Learning Analytics amp; Knowledge Conference, Association for Computing Machinery, New York, NY, USA. p. 576–577. doi:10.1145/3027385.3029467.

[84] White, R.W., Dumais, S.T., Teevan, J., 2009. Characterizing the influence of domain expertise on web search behavior, in: Proceedings of the Second ACM International Conference on Web Search and Data Mining - WSDM '09, ACM Press, New York, New York, USA. p. 132. doi:10.1145/1498759.1498819.

[85] Whitehill, J., Mohan, K., Seaton, D., Rosen, Y., Tingley, D., 2017. Mooc dropout prediction: How to measure accuracy?, in: Proceedings of the Fourth (2017) ACM Conference on Learning @ Scale, Association for Computing Machinery, New York, NY, USA. p. 161–164. doi:10.1145/3051457.3053974.

[86] Wilkowski, J., Deutsch, A., Russell, D.M., 2014. Student skill and goal achievement in the mapping with google mooc, in: Proceedings of the First ACM Conference on Learning @ Scale Conference, Association for Computing Machinery, New York, NY, USA. p. 3–10. doi:10.1145/2556325.2566240.

[87] Xiao, J., Jiang, B., Xu, Z., Wang, M., 2015. The usability research of learning resource design for MOOCs, in: Proceedings of IEEE International Conference on Teaching, Assessment and Learning for Engineering: Learning for the Future Now, TALE 2014, Institute of Electrical and Electronics Engineers Inc.. pp. 277–282. doi:10.1109/TALE.2014.7062640.

[88] Yi, X., Hong, L., Zhong, E., Liu, N.N., Rajan, S., 2014. Beyond clicks: Dwell time for personalization, in: Proceedings of the 8th ACM Conference on Recommender Systems, ACM, New York, NY, USA. pp. 113–120. doi:10.1145/2645710.2645724.

[89] Yu, R., Gadiraju, U., Holtz, P., Rokicki, M., Kemkes, P., Dietze, S., 2018. Predicting user knowledge gain in informational search sessions, in: 41st International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2018, Association for Computing Machinery, Inc. pp. 75–84. doi:10.1145/3209978.3210064.

[90] Zhang, D., Zhou, L., Briggs, R., Nunamaker, J., 2006. Instructional video in e-learning: Assessing the impact of interactive video on learning effectiveness. Information Management 43, 15–27. doi:10.1016/j.im.2005.01.004.

[91] Zhang, N., Biswas, G., Dong, Y., 2017. Characterizing students' learning behaviors using unsupervised learning methods, pp. 430–441. doi:10.1007/978-3-319-61425-0_36.