# Return of the Student: Predicting Re-Engagement in Mobile Learning

Maximillian Chen
Cornell University
mlc294@cornell.edu

René F. Kizilcec
Cornell University
kizilcec@cornell.edu

## ABSTRACT

Mobile learning platforms cater to intermittent microlearning by lowering the barrier for re-engaging in the learning process after a period of disengagement. We examine student re-engagement in the context of an SMS-based mobile learning platform, how to predict it and how it differs from disengagement. In a sample of $87,651$ Kenyan students, we analyze data on $1,196,780$ quiz attempts, finding that 36.3% of students who disengage for a week or more eventually re-engage on the platform. They spend more time on quizzes early on than students who stay disengaged. A Random Forest classifier trained on two days of student activity logs predicts disengagement and re-engagement with similar performance: F1 scores of 81.2% and 80.9%, respectively. The prevalence of re-engagement in mobile learning calls for more research into this behavioral outcome.

## 1. INTRODUCTION

As the world becomes increasingly connected through mobile technology, mobile devices are becoming an increasingly viable medium for education. Not only are mobile phones more affordable than traditional personal computers, but mobile devices have shallower learning curves, as they require lower levels of literacy and training [14]. The accessibility of mobile technology is especially advantageous in resource-constrained areas. It provides students access to educational resources without having to make substantial economic trade-offs associated with desktop computers. Given the rapid development of mobile computing power, many people in developing economies are predicted to skip purchasing desktop computers altogether and instead adopt mobile devices [6]. In comparison to traditional online learning platforms, mobile learning platforms remain relatively understudied despite their promise for accessibility.

A common concern with self-directed learning tools is that students do not stay engaged on the learning platform for long. The issue of disengagement, defined as a drop in student activity on the platform, has been studied extensively, for instance in the context of Massive Open Online Courses (e.g. [16, 15, 12, 8]). However, it remains largely unstudied in the context of mobile learning environments. Student engagement patterns likely vary between desktop and mobile learning environments, considering how many different applications are available [2] and how deeply embedded mobile devices are in people's everyday lives. In fact, mobile learning platforms have been found to provide unique opportunities for microlearning sessions, where learning tasks are broken into shorter chunks that can be managed "on-the-go" [4]. Especially considering the low barriers to entry and exit in most mobile learning applications, it is unsurprising that a sizable proportion of students engage and disengage freely, which can result in longer gaps of inactivity. These intermittent usage patterns require that we consider re-engagement as a distinct behavior in mobile learning and how it compares to disengagement. Insights from this work can advance our understanding of how mobile learning works in practice and how platforms may support at-risk students through intervention.

In this research, we propose definitions of disengagement and re-engagement in mobile learning, analyze differences in behavior between disengaging and re-engaging students, and apply supervised machine learning approaches to predicting disengagement and re-engagement in mobile learning. We find that 36.36% of students who disengage for a week eventually re-engage on the platform within two weeks. A Random Forest classifier trained on two days of student log data can predict re-engagement after two weeks with an F1 score of 80.9%, showing that early platform activity is indicative of which students will return later on.

## 2. BACKGROUND

### 2.1 Beyond Student Disengagement

Before defining re-engagement, we need to formally define its prerequisite: disengagement. Defining disengagement in mobile learning platforms presents a challenge, because many such platforms are inherently less rigid and prescriptive in their learning design compared to online learning environments such as massive open online courses (MOOCs). MOOCs tend to lay out a clear path through course materials with deadlines, while many mobile learning platforms provide more room for self-directed learning and agency in choosing a learning path. This calls for an updated definition of disengagement for the context of mobile learning.

Disengagement is defined conceptually as a "lack of engage-

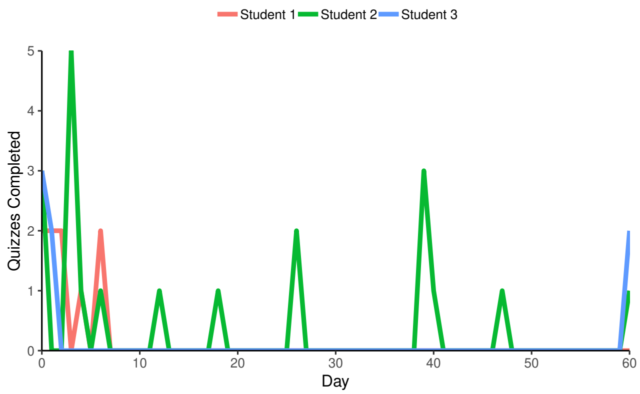


**Figure 1: Three students' daily activity for 60 days after signing up: student 1 disengages; student 2 stays engaged; student 3 disengages and re-engages.**

ment," and has been operationalized in terms of students' interaction with or completion of learning objectives, depending on the structure of the learning platform [12]. In past studies in the context of MOOCs, disengagement has been defined as a "lack of interaction" [9, 1], the point in time where a student "fails to submit any further assignments" [15], failure to earn a course certificate [5], failure to complete a set of modules [3] or a lack of platform interaction combined with a lack of progress towards course completion [12]. Despite the many studies defining and predicting engagement in MOOCs, research on modeling engagement in mobile learning is scarce [7] and definitions of disengagement may not translate well from MOOCs to mobile learning. Definitions focusing on course completion do not apply in a context that is unlike a course, and definitions focusing on an absence of platform interaction may incorrectly label students who disengage early but return to the platform later on. Mobile learning platforms offer more opportunities for "microlearning" sessions in which students are able to learn in short bursts sporadically or "on-the-go" [4].

A recent study on engagement patterns in mobile learning on the same platform as in this research found students tend to be engaged in learning activities during the first few days after signing up but disengage shortly thereafter. In fact, 75% of all students disengaged within two days of registering, and even among the cluster of engaged students, 68% of them appeared to disengage in the first ten days [7]. However, whether a students has completely disengaged can be unclear at first sight. Consider the three actual students whose activity over time is visualized in Figure 1. All three are engaged in the first week, but student 1 disengages and never re-engages, while student 2 is inactive during the second week but occasionally returns to complete quizzes over the next two months. Student 3 was engaged on the day of registration but then disengaged for 60 days before re-engaging. We therefore define a re-engaging student as one who disengages but then returns to the mobile learning platform. This more accurately characterizes student behavior in the long run and with some additional granularity.

The ability to distinguish re-engaging students from disengaging students has practical applications, such as for an automated student support system. The system could send different kinds of text messages or notifications to students who are classified as disengaging (i.e. not ever re-engaging) based on their activity in the first few days. By targeting students based on their predicted behavior, providers can tailor reminders to groups of students to highlight learning opportunities without alienating students who are already likely to re-engage in the absence of nudging.

## 2.2 Predicting Student Engagement

As there have not been any large-scale studies predicting student engagement in mobile learning to date, we build on a large literature on predicting student engagement and intervention systems in the context of MOOCs [16, 12, 15]. As is the case in MOOCs, a vast majority of students on mobile learning platforms eventually disengage. Any supervised learning approach in which labels correspond to engagement/disengagement would therefore suffer from class imbalance, i.e., the distribution of class labels is heavily skewed [10]. A naive classifier could simply predict the majority label for all instances and achieve a high degree of accuracy without successfully identifying actual engagement. Nagrecha and colleagues [12] addressed this issue by re-sampling their training data to balance the distribution of their labels, as was done in prior work predicting student disengagement [11]. Due to class imbalance, model accuracy can be a misleading evaluation metric, and prediction recall is frequently used as a substitute. Likewise, in this study, we face a heavy imbalance in class labels (very few students re-engage). We therefore opt to re-sample our data during training and evaluate our models using both recall and F1 score.

The user interface of mobile learning platforms tends to be simpler than ones designed for larger computer screens. For example, MOOCs tend to have more advanced platform features than mobile learning applications, such as video playback options and non-traditional assessment types. Prior studies have focused on engineering features relevant to interaction with video lectures, such as "number of straight-through video plays" or "number of video views per session" [9]. But clickstream data (interaction logs) are more informative about interactions with the structure of a platform than any specific course, which is why they were found to offer strong predictive power when analyzing data across multiple courses on a learning platform [17]. We pose two research questions in this study:

**RQ1.** How does the behavior of re-engaging students compare with those who stay disengaged?

**RQ2.** What features are predictive of student re-engagement?

# 3. METHODS

## 3.1 Platform Background & Dataset

We study re-engagement on a text message-based mobile learning platform called Shupavu 291. It has been used by over 5 million students and it offers content for over 800 distinct curricula. The platform was developed by Eneza Education[1] to provide a learning resource in regions with

[1]https://enezaeducation.com/

Table 1: **Daily student activity features for two days.**

| Feature Name | Definition |
|---|---|
| time.i | Time spent on day i |
| nlessonsfinished.i | Num. of lessons completed on day i |
| nask.i | Num. of questions asked on day i |
| n_quizzes.i | Num. of quizzes completed on day i |
| avg_solve_time.i | Avg. time to complete quizzes on day i |
| n_unique_quizzes.i | Num. of unique quizzes completed on day i |
| nsummary.i | Num. of quiz results viewed on day i |
| nhw_tools.i | Num. homework tools (e.g. dictionary) used on day i |

limited access to education. Shupavu 291's user base primarily consists of Kenyan students, though its influence is growing in other African countries. The platform was designed by a group of Kenyan teachers, and the course materials align with the topics and learning outcomes of the Kenyan national curriculum for numerous subjects in primary and secondary education. Every interaction with Shupavu 291 is via text message. Students navigate through menus and quizzes by sending a text message containing a number corresponding to a menu item from the options relayed to them. Students are able to choose from a variety of grade-specific subjects such as "Fractions" and "Kiswahili." For a given subject, students choose a specific topic and receive compact lecture notes followed by a quiz (generally five multiple-choice questions). Quiz questions follow the menu format and are sent individually; students receive instant feedback on correctness along with an explanation. Students may retake quizzes as many times as they like or use the "Ask-A-Teacher" feature to ask teachers for help.

Shupavu 291 stores a record for every quiz or platform interaction a student completes. The dataset used consists of 21,302,582 platform actions, including 1,196,780 quiz attempts, from 87,651 students in Kenya. Data beyond self-reported grade level and platform interactions for each student is completely de-identified. For the purpose of this research, we construct two sub-samples, where each one is used to solve a separate prediction problem. The first sample consists of the 87,651 students who completed at least one quiz on Shupavu 291 (an indicator of their willingness to engage with content). The second sample consists of those 63,120 students in the first sample who exhibited a seven-day period of inactivity (i.e. disengagement). The sample definitions are explained further in the next section.

## 3.2 Defining the Prediction Task & Features

We define two separate prediction tasks: predicting disengagement, and predicting eventual re-engagement. A disengaging student is defined as one who is inactive (here, not attempting quizzes) for at least seven consecutive days. A re-engaging student is defined as one who has disengaged and then is active (here, attempts quizzes) for at least two different days within the 14-day period following the period of inactivity. As in most disengagement prediction problems [12], we found a significant imbalance in observed labels for both disengagement and re-engagement: 72.01% of students were labeled as disengaging, and 63.68% of them were labeled as remaining disengaged (i.e. not re-engaging). We thus trained our classifiers on data that was randomly re-sampled to achieve a more balanced label distribution.

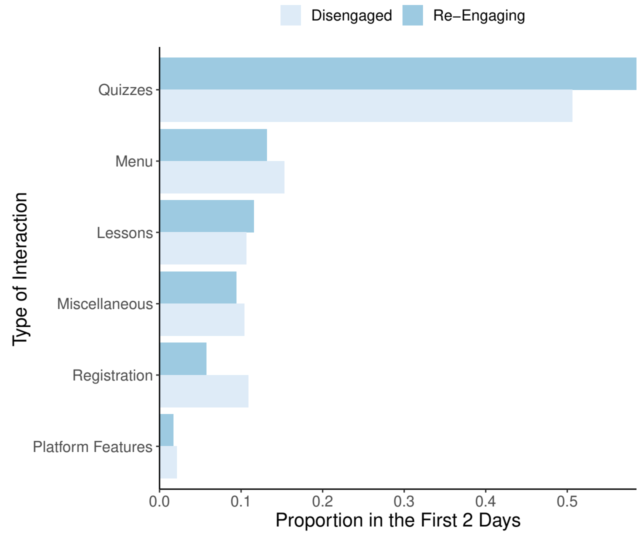


**Figure 2: Distribution of the number of interactions with different parts of the platform in the first two days for students who re-engage and those who stay disengaged.**

Due to the rapid decline in student engagement after registration, we devise features to capture activity on each student's first two days on the platform. We expect early engagement to be predictive of disengagement and re-engagement. The features, defined in Table 1, capture how students interact with key components of the Shupavu 291 application and generalize across multiple subject areas, similar to the method used by Taylor and colleagues [15]. Min-max normalization is performed for each feature for each student's first two days. We fit a Random Forest (RF) to predict whether students disengage, then another RF to predict their re-engagement. Model performance is evaluated using Recall (as suggested in [12]) and F1 scores (the harmonic mean of Precision and Recall), where a "true positive" is a student who disengages (or remains disengaged). We optimize model hyper-parameters to maximize F1 scores through exhaustive 5-fold cross-validated grid search using scikit-learn [13].

# 4. FINDINGS

We find that more than a third (36%) of students who disengage (seven days of inactivity) eventually re-engage on the Shupavu 291 mobile learning platform. The prevalence of re-engagement in this learning context speaks to the importance of considering this engagement pattern in mobile learning more broadly. To address the first RQ about differences between disengaging and re-engaging students, we compare student activity in the first two days after registering on the platform. Actions on Shupavu 291 are grouped into six categories:

- *Registration:* managing Shupavu 291 subscriptions.
- *Menu:* navigating the menu structure.
- *Lessons:* using course material, e.g. completing lessons.
- *Quizzes:* answering quiz questions, checking quiz grades, or starting quizzes.
- *Platform Features:* using Shupavu 291-specific resources,

e.g. the dictionary or ask-a-teacher feature.
- *Miscellaneous:* any other interaction, e.g. promotional events and features.

Overall, disengaging and re-engaging students behave similarly, spending most of their time interacting with quizzes (Figure 2). However, re-engaging students interact significantly more with quizzes (56.00% v. 47.57%, $\chi^2 = 25083, p < 0.001$) and slightly more with lessons (11.45% v. 10.34%, $\chi^2 = 1109.6, p < 0.001$, while disengaging students have more registration events (13.30% v. 7.25%, $\chi^2 = 35827, p < 0.001$). Having a greater proportion of registration events may be an indication that students who stay disengaged were already spending less time engaging with Shupavu 291 even within their first two days. A greater proportion of academic (quiz and lesson) events is likely an indication that students who eventually re-engage were more active students early on. The finding that re-engaging students engage with more academic events early on is notable, as quizzes and lessons are the core functions of Shupavu 291.

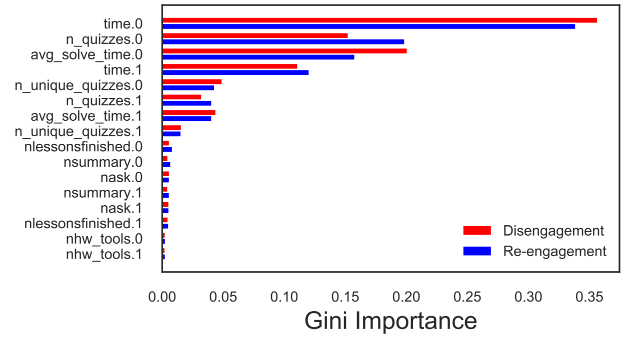


**Figure 3: RF Gini Importance by prediction problem. Number of quizzes completed and time spent after the first day are more important predictors of re-engagement than disengagement.**

## 4.1 Predicting Modes of Engagement

We fit an RF[2] to predict disengagement and re-engagement using a set of features that capture early platform activity (Table 1). The model achieved good results for the disengagement prediction task, with a testing F1 score of 81.21% and Recall of 83.06%. Fitting the same RF to predict re-engagement received comparable performance: 80.91% F1 score and 84.19% Recall. This suggests that it is possible to train a useful classifier for both behaviors using early engagement features.

To better understand which kinds of early behaviors predict each outcome, we compare variable importance scores between the models in Figure 3. The number of quizzes completed on both day 0 and day 1, time spent on day 1, and number of platform features (questions asked, homework tools, quiz summaries) used on both day 0 and day 1 are more important for predicting re-engagement, whereas the other features are more important for predicting disengagement. This suggests that quiz engagement and diversity of platform usage is especially predictive of a student's likelihood to re-engage, though many of these characteristics are also predictive of disengagement. The importance of time spent on day 1 for predicting re-engagement is notable because it indicates that long-term behavior is related to sustained activity. Aside from the finding that diversity of platform usage is more important for predicting re-engagement, Figure 3 also suggests that specific platform feature usage (e.g. "Ask A Teacher") is not as indicative of student engagement as in prior work with MOOCs [9]. Overall, we find that early usage behavior is predictive of students' subsequent engagement pattern, which provides a basis for developing automatic interventions to better support students.

# 5. DISCUSSION

This study shows the prevalence of re-engagement in mobile learning. This behavioral outcome can be defined in many different ways and the optimal choice will depend on the context of the learning environment and broader goals of the predictive model. In particular, the periods of time and the thresholds of activity to determine dis- and re-engagement can be tweaked to fit context and goals. In the context of Shupavu 291, which has rapid disengagement, most periods of inactivity occur soon after registration. Most predictive models are not also explanatory models and this is no exception. While it is feasible to predict how a student will behave, it is unclear why they (choose to) behave in that way. A student who is active early on but disengages for a seven-day period several weeks after registering could be treated differently than one who disengages early on for the purpose of targeted intervention. Yet more work is needed to discern how to support students differentially in light of their predicted outcome.

One of the limitations of this study lies in how the second prediction problem is set up. The model is trained only on the subset of students who disengage, because by our definition, a student who does not disengage cannot re-engage. Alternatively, we could have taken the output from the disengagement prediction model and predicted the joint likelihood of disengaging and re-engaging. However, this would have introduced a great deal of uncertainty from the disengagement task into the re-engagement task. Another alternative is to set up the re-engagement prediction task as identifying students who disengage and then re-engage; however, in this case, all other students are then a mix of those who disengage completely and those who remain engaged the entire time—two groups which exhibit very different behaviors. Restricting the sample to only disengaged students gives up some information, but provides a clear basis of comparison for predicting disengaging and re-engaging students.

This research contributes an empirical treatment of student re-engagement in mobile learning and one of the first large-scale studies of student interaction with a mobile learning platform, especially in the developmental context of Sub-Saharan Africa, where mobile learning provides students with affordable access to study resources outside of formal schooling. We find it is possible to predict and distinguish between disengaging and re-engaging students using early clickstream data, providing a foundation for more research into patterns of re-engagement.

[2] $1,000$ trees, 2 samples/split min., 1 sample/leaf min., 25 tree depth limit, *Gini* criterion, optimized for F1 score

# 6. REFERENCES


[1] G. Balakrishnan and D. Coetzee. *Predicting student retention in massive open online courses using hidden markov models*. Technical Report No. UCB/EECS-2013-109, 2013.

[2] I. Blair. *Mobile App Download and Usage Statistics (2019)*, Feb 2019. https://buildfire.com/app-statistics/.

[3] J. Dillon, N. Bosch, M. Chetlur, N. Wanigasekara, G. A. Ambrose, B. Sengupta, and S. K. D'Mello. Student emotion, co-occurrence, and dropout in a mooc context. In *Proceedings of the 9th International Conference on Educational Data Mining*, 2016.

[4] T. Dingler, D. Weber, M. Pielot, J. Cooper, C.-C. Chang, and N. Henze. Language learning on-the-go: opportune moments and design of mobile microlearning sessions. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*, pages 1–12, 2017.

[5] J. He, J. Bailey, B. I. Rubinstein, and R. Zhang. Identifying at-risk students in massive open online courses. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

[6] G. Jones. *Clickology: What Works in Online Shopping and How Your Business Can Use Consumer Psychology to Succeed*. Hachette UK, 2013.

[7] R. F. Kizilcec and M. Chen. (in press). Student engagement in mobile learning via text message. In *Proceedings of the Seventh (2020) ACM Conference on Learning @ Scale*, 2020.

[8] R. F. Kizilcec, C. Piech, and E. Schneider. Deconstructing disengagement: analyzing learner subpopulations in massive open online courses. In *Proceedings of the third international conference on learning analytics and knowledge*, pages 170–179, 2013.

[9] M. Kloft, F. Stiehler, Z. Zheng, and N. Pinkwart. Predicting mooc dropout over weeks using machine learning methods. In *Proceedings of the EMNLP 2014 workshop on analysis of large scale social interaction in MOOCs*, pages 60–65, 2014.

[10] R. Longadge and S. Dongre. Class imbalance problem in data mining review. *arXiv preprint arXiv:1305.1707*, 2013.

[11] C. Márquez-Vera, A. Cano, C. Romero, A. Y. M. Noaman, H. Mousa Fardoun, and S. Ventura. Early dropout prediction using data mining: a case study with high school students. *Expert Systems*, 33(1):107–124, 2016.

[12] S. Nagrecha, J. Z. Dillon, and N. V. Chawla. Mooc dropout prediction: lessons learned from making pipelines interpretable. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 351–359, 2017.

[13] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[14] I. Quadir. Form, transform, platform: How the ubiquity of mobile phones is unleashing an entrepreneurial revolution. *Innovations: Technology, Governance, Globalization*, 7(4):3–12, 2012.

[15] C. Taylor, K. Veeramachaneni, and U.-M. O'Reilly. Likely to stop? predicting stopout in massive open online courses. *arXiv preprint arXiv:1408.3382*, 2014.

[16] J. Whitehill, K. Mohan, D. Seaton, Y. Rosen, and D. Tingley. Delving deeper into mooc student dropout prediction. *arXiv preprint arXiv:1702.06404*, 2017.

[17] J. Whitehill, J. Williams, G. Lopez, C. Coleman, and J. Reich. Beyond prediction: First steps toward automatic intervention in mooc student stopout. *Available at SSRN 2611750*, 2015.