Tools for Predicting Drop-off in Large Online Classes

Justin Cheng,
Chinmay Kulkarni,
Scott Klemmer
353 Serra Mall
Stanford University
Stanford CA 94305 USA

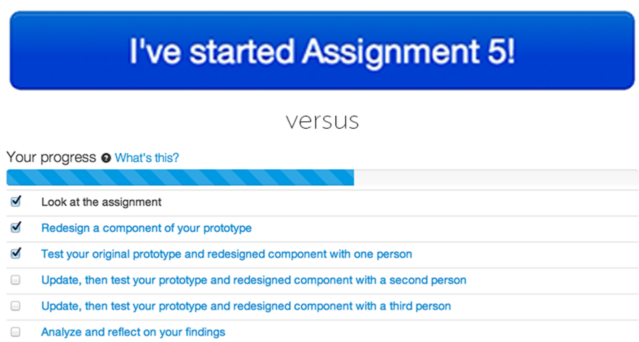

Figure 1: Screenshots of experimental prototypes. **Above** The single "started" button, **Below** The action plan prototype

CSCW'13 Companion, Feb. 23–27, 2013, San Antonio, Texas.
ACM 978-1-4503-1332-2/13/02.

Abstract
This paper describes two diagnostic tools to predict students are at risk of dropping out from an online class. While thousands of students have been attracted to large online classes, keeping them motivated has been challenging. Experiments on a large, online HCI class suggest that the tools these paper introduces can help identify students who will not complete assignments, with an F_1 score of 0.46 and 0.73 three days before the assignment due date.

Author Keywords
Online classroom interventions; micro-commitment; action plans; online education

ACM Classification Keywords
H.5.2 [**Information interfaces and presentation (e.g., HCI)**]: User Interfaces.

Introduction
While online classes have allowed instructors to attract and reach a large, diverse population of students, a significant proportion of students who sign up fail to complete these classes.

Of the 45 thousand students who enrolled in the Spring 2012 offering of Stanford University's online HCI class,

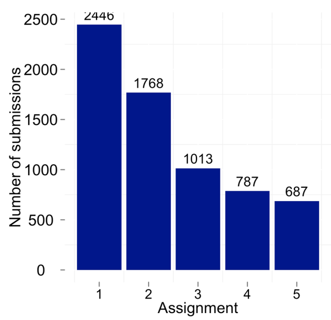

Figure 2: Student retention for the Spring 2012 edition of the class. Fewer students submit assignments later in the class.

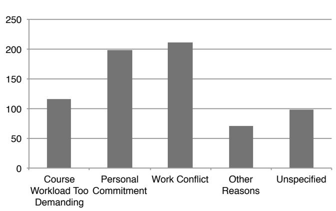

Figure 3: Student reasons for disengagement: Spring 2012. A large fraction of students dropped out because of work conflicts and personal commitments.

only 2400 (5%) complete the first design assignment. This low percentage is possibly because first assignment requires several hours of work to complete, and many students sign up only to preview the class material.

However, even highly motivated students who complete the first assignment don't stay motivated for long: a mere 28% complete the class and turn in the final assignment (Figure 3). These results mirror the low retention rates (anecdotally, 25% to 35%) typically seen in many other online classes.

In contrast, the average retention rate for college freshman programs in the US is 75 percent [1]. While college students may have stronger extrinsic motivations for completing courses (ex. parental pressure or obtaining a degree for future employment), a survey we conducted indicates that it is possible to design interventions to online class retention.

To understand why students dropped out, we invited students from the Spring 2012 edition of the class to a voluntary end-of-class survey. The survey asked students who did not watch any videos or complete any assignments in the last 3 weeks of the class about their reasons for dropping out (Figure 3). Students cited work conflict (30.3%) or personal commitment (28%) as primary reasons for dropping out, suggesting that students may benefit from help in prioritizing the class.

It is difficult to personally identify students who need assistance in online classes given the low instructor to student ratio. This paper introduces two scalable tools – a micro-commitment button, and an action plan – to help identify students at risk of dropping out of a class. We tested these tools in the Fall 2012 edition of the online HCI class, and report results on this field deployment.

Related work

The concept behind the micro-commitment button is inspired by previous research which found that people construct consistent narratives about their actions [4]. To that effect, commitment contracts, or requiring students to sign a learning contract is positively correlated with better performance and attitudes [5]. As we only require students to click on a button to indicate that they've started on an assignment, we feel that this is a weaker form of commitment than a commitment contract, or "micro-commitment".

Action plans, or checklists, have also been shown to be effective in helping participants recall and complete tasks [2, 3, 7]. Such plans are especially useful in creative work, where people tend to procrastinate because of the open-ended nature of such tasks [8]. Prior work suggests that providing concrete, personalized action plans helps increase task completion rates [3].

System design

We designed two different prototypes to identify students who would potentially not submit assignments. Both are based on the hypothesis that students' interaction with the assignment is indicative of their motivation. Therefore, both prototypes elicit information from students about their progress through the assignment. The design challenge is eliciting signals from students early enough to allow instructors to target them specially.

The first prototype shows participants a single button with the text "I've started on this Assignment", which turns green when they click on it (Figure 1). We display this button at the top of the assignment page.

The second prototype is a multi-step action plan based on the assignment instructions. Action plans specify concrete steps required to complete a task. Similar to the micro-commitment button, the action plan appears at the top of the assignment web page in the form of a checklist, and specific steps also appear at relevant points in the assignment's description. A student can check off steps as they complete the assignment, and a progress bar displays progress through the assignment. The "zeroth" step in each action plan is "View this assignment", which is checked off automatically when the student first visits the assignment page.

Experiment

Participants and design

To test our prototypes, we conducted a between-subjects experiment amongst students in the online HCI class. In all, 1207 students participated in this experiment. Participation was uncompensated and voluntary. Participants were chosen amongst students who submitted the second assignment in the class, and therefore had put in about 10 hours of work into the class. The experiment had three conditions: a micro-commitment condition ($n = 248$), an action-plan condition ($n = 720$) and a wait-list control condition ($n = 239$).

Participants in the wait-list condition were told that we were slowly rolling out experimental features to the entire class, and that they had been put on a waiting-list. Using a wait-list control has similar benefits to using a placebo condition [9].

Procedure

Participants saw the prototype manipulation one week before the next assignment is due. Participants in the control condition saw the standard course website with no changes. Below, we predict engagement based on whether or not a student had taken an action by a given day.

Results

Metrics for measuring accuracy

Both prototypes require students to take an explicit action (clicking a button/checkbox) before we can predict that they are engaged (students who don't click are assumed to be disengaged). However, students might take this action at different times in the week. Therefore, it is important to not only know how accurately prototypes can predict engagement, but how *early* they can do so. Therefore, we report three measures– recall (the percentage of engaged students we can detect), precision (the pecentage of students that are actually engaged amongst those that we detect as engaged.) and F_1 score, a harmonic mean of precision and recall.

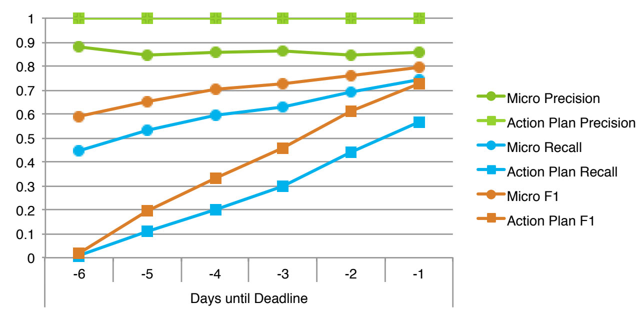

Figure 4: Precision, recall and F1 score for both prototypes

Both started button and action plan effective

Overall, both the started button and the action plan are effective at detecting engaged students. Figure 4 shows Precision, recall and F1 numbers for both prototypes. For the action plan, we only analyze results for the first step that students need to explicitly take action on (the zeroth step was automatically checked). We also disregard when students completed subsequent steps, to make the comparison with the started button fairer (which only can elicit one bit of information). Three days before the assignment is due, F_1 numbers for the started button and the action plan are 0.73 and 0.46 respectively.

Action plans more accurate, started button an earlier signal
The action plan prototype has very good precision through out, but its recall is only high a day before the assignment is due. This suggests that even students who are engaged with the class don't check off their first step until a few days before the deadline. In contrast, the started button has lower precision than the action plan, but has much higher recall.

Action plans increase student engagement
Using a mixed linear model reveals that assignment completion is significantly higher when action plans are presented to the user ($p < 0.05$). Approximately, 45% of students in all conditions submitted assignments.

Conclusions and future work

Going beyond detection
This poster explores the use of the action plan and micro-commitment button as a tool to detect student engagement. In the future, we plan to use these signals to improve engagement. In particular, these two tools that we propose can provide a strong baseline for machine learning algorithms that aim to predict students at risk of dropping out. Online classes have a potential to teach large numbers of students distributed around the world (Coursera alone estimates that more than two million students are enrolled on its platform [6]). Keeping these students engaged will be important for advancing this new medium of learning.

Acknowledgements
We thank our colleagues at Coursera for helping us build the experimental system. This work is supported by a Hasso Plattner Design Thinking Program grant and an SAP Stanford Graduate Fellowship.

References

[1] US News Survey Results. http://hci.st/usnews.

[2] Allen, D. *Getting Things Done. The Art of Stress-Free Productivity*. Penguin, January 2003.

[3] Bellotti, V. e. a. What a to-do: studies of task management towards the design of a personal task list manager. In *Proc. SIGCHI '04*,

[4] Cialdini, R. Science and practice.

[5] Kasworm, C. E. An examination of self-directed contract learning as an instructional strategy. *Innovative Higher Education 8* (1983), 45–54.

[6] Levin, T. Education site expands slate of universities and courses. http://hci.st/courseraMillion, 2012.

[7] Luszczynska, A. An implementation intentions intervention, the use of a planning strategy, and physical activity after myocardial infarction. *Social Science and Medicine* (2006).

[8] O'Donoghue, T., and Rabin, M. Choice and Procrastination. *Quarterly Journal of Economics 116*, 1 (2001), 121–160.

[9] Speca M., e. a. A randomized, wait-list controlled clinical trial: the effect of a mbsr program on mood and symptoms of stress in cancer outpatients.