

Supporting the Encouragement of Forum Participation

Aashna Garg
Stanford University
aashna94@stanford.edu

Andreas Paepcke
Stanford University
paepcke@cs.stanford.edu

1. INTRODUCTION

Since 2011 many courses in universities have used the Piazza forum facility for their internal courses. As [4] and others have shown, forum participation can be beneficial for all students. Are there potential intervention points in time during an academic period when encouragements (e.g. [6]) might strengthen forum use?

We analyze the forum contribution rates during forty offerings of 12 college courses, reaching back to 2011. By ‘offering’ we denote a course taught during a particular quarter. Multiple offerings of the same course afforded us longitudinal observations. We looked for significant leaps in forum contribution, hypothesizing that those times are important intervention points.

From a social graph model of forum activity we computed weighted out degrees and page ranks for each student. The out degrees reflect the number of posts a student contributes. We performed change point analyses through bootstrap procedures over the CUSUM data of the post contributions through each quarter. We thereby identified significant week by week changes in the rate at which the top ten percent of forum contributors post messages. We hypothesize that such change points might be appropriate encouragement opportunities, and we find that sudden rate shifts do occur along, sometimes in regular patterns, primarily in science and engineering courses. We propose and demonstrate the use of control charts to monitor forum traffic, and show how the historic data can be used to provide personalized encouragement messages.

2. ANALYSIS

Our question was “when would be good times during a course for encouraging students who lag behind in forum contributions?” We hypothesize that times when the top 10% of contributors speed up, other students might be encouraged to do the same.

The results indicate that week six is particularly likely to experience changes in posting rates. The studied university operates on a quarters schedule, we can hypothesize that the traffic is related to midterms. Week eight might be related to final projects coming due not too far out. Note, however, that the number of offerings (shown in parentheses with the course names) differ across courses. Thus 20% for urban studies means only one of five offerings experienced a significant posting acceleration. The only courses with an

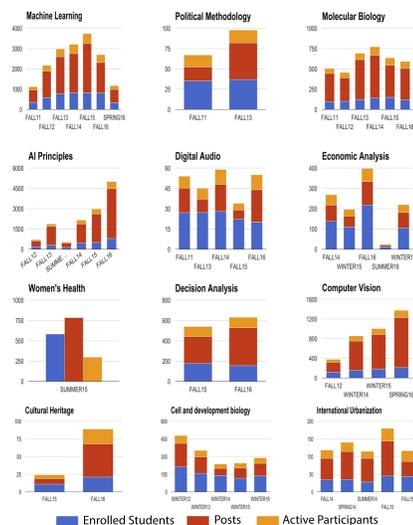


Figure 1: Number of students enrolled in the course, number of students active in the forum, and number postings.

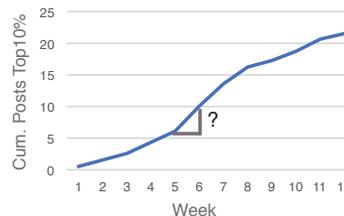


Figure 2: Cumulative number of postings by top 10% contributors throughout an academic quarter (machine learning class).

appreciable number of change points among their offerings are engineering courses. The humanities and social sciences, while using forum facilities, have not included the forum as a central discussion hub. The numbers of students attending these courses are also smaller than the science/engineering classes.

The forum change point computations we outlined above require data from all weeks of an offering to be available. In the presence of historic data this requirement is not a problem. But what could an instructor do to discover unusual posting frequencies while the offering is running?

2.1 Control Charts for Forum Alerts

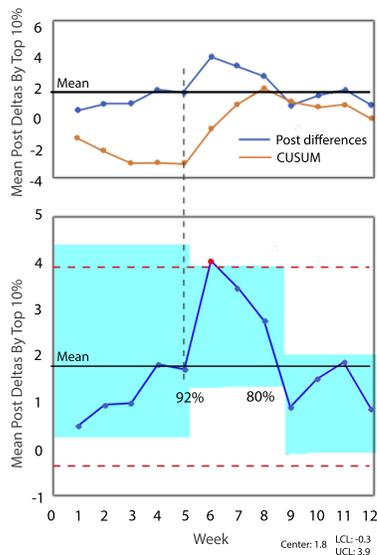


Figure 3: Change points in forum posting time series (machine learning class).

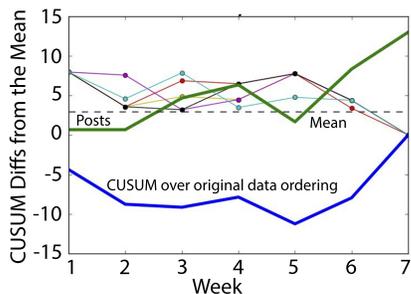


Figure 4: CUSUMs from permuted weekly post contribution rates hover around the mean, green line indicating the number of posts.

Figure 3's lower chart hints at a possibility we have not yet discussed. The horizontal dotted lines are *process limits*, a term from process control practice [3]. The limits bound the values within which a process is expected to vary. For industrial processes the variations might lie around a known optimal operating level, such as a temperature. When no such level is known from the domain, a mean can be used. Our process limits denote $2 * \sigma$ distance from the mean.

In the context of instructional forum use it would be possible to detect points, such as the one above the upper control limit in Figure 3. Such change to above-normal might indicate confusion among the students, or the discovery of an exciting topic. Either way, the instructor's attention might be warranted, as might be pointing passive students to the increased activity.

2.2 Personalized, Quantitative Encouragement

The data framework we discussed could also be deployed to provide personalized encouragement for passive students. Rather than admonishing students for past passivity, a forward-looking nudge could be provided. Figure 5 illustrates this option. The dot represents one student and their contribution as of week six: two postings. Two of many possible options are shown in the Figure for catching up to the top-

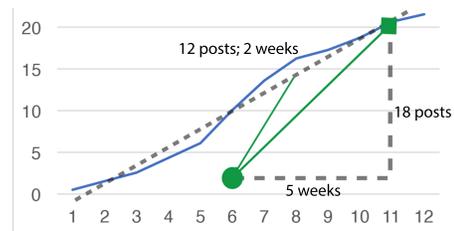


Figure 5: Personalized, quantitative nudges: catching up with the top-10% within five weeks, or two weeks.

10% contributors. These, and other options are derived by drawing a line from the student's position to an intersection with the regression line. That line is known from past course offerings. The larger the slope of the connecting line, the more aggressive the plan for catching up. For example, the long line would catch the student up within five weeks, assuming a weekly rate of $(20 - 2)/5 \approx 4$ messages per week.

Alternatively, the shorter line would call for six messages per week, to catch up within two weeks. The square at the end of the long option is intended to represent a slider that the student could run along the regression line to make a plan. The number of required weekly messages would be updated continuously as the student operates the slider. While this sketch is not the intended user interface, it illustrates the ideas of using past and current forum data to provide (i) forward-looking encouragement in place of recrimination, and (ii) to empower the student by personalizing the message, and providing a tool for planning. Studies are needed to determine whether postings in response to even such personalized messages prove beneficial.

For discussions on the impact of forum participation, see [6, 2, 5], and [1].

3. CONCLUSION AND FUTURE WORK

The most important next step is to test whether existing change points truly are effective encouragement moments. We also plan to use page rank measurements to see TA work sharing patterns.

4. REFERENCES

- [1] A. Anderson, D. P. Huttenlocher, J. M. Kleinberg, and J. Leskovec. Engaging with massive online courses. *CoRR*, abs/1403.3100, 2014.
- [2] C. G. Brinton, M. Chiang, S. Jain, H. Lam, Z. Liu, and F. M. F. Wong. Learning about social learning in moocs: From statistical analysis to generative model. *CoRR*, abs/1312.2159, 2013.
- [3] Department of Commerce. Nist/sematech e-handbook of statistical methods. Electronic book, 2012.
- [4] J. Huang, A. Dasgupta, A. Ghosh, J. Manning, and M. Sanders. Superposter behavior in mooc forums. In *Proceedings of the First ACM Conference on Learning @ Scale Conference, L@S '14*, pages 117–126, New York, NY, USA, 2014. ACM.
- [5] D. Yang, M. Wen, and C. Rose. Peer influence on attrition in massive open online courses, 2014.
- [6] M. Yeomans and J. Reich. Planning prompts increase and forecast course completion in massive open online courses. In *Proceedings of the Seventh International Learning Analytics & Knowledge Conference, LAK '17*, pages 464–473, New York, NY, USA, 2017. ACM.