LMS Course Design As Learning Analytics Variable

John Fritz
Univ. of Maryland, Baltimore County
1000 Hilltop Circle
Baltimore, MD 21250
410.455.6596
fritz@umbc.edu

ABSTRACT

In this paper, I describe a plausible approach to operationalizing existing definitions of learning management system (LMS) course design from the research literature, to better understand instructor impact on student engagement and academic performance. I share statistical findings using such an approach in academic year 2013-14; discuss related issues and opportunities around faculty development; and describe next steps including identifying and reverse engineering effective course redesign practices, which may be one of the most scalable forms of analytics-based interventions an institution can pursue.

Categories and Subject Descriptors

K.3.1 [**Computer Uses in Education**]: Collaborative learning; Computer-assisted instruction (CAI); Computer-managed instruction (CMI); Distance learning.

General Terms

Algorithms, Design, Human Factors, Measurement, Standardization, Theory

Keywords

Course Design, Instructor Pedagogy, Learning Analytics Methodology

1. INTRODUCTION OF PROBLEM

Given wide spread use of the learning management system (LMS) in higher education, it is not surprising this form of instructional technology has frequently been the object of learning analytics studies [1, 2, 3, 4, 5]. While methods and results have been mixed in terms of predicting student success, let alone leading to actual, effective and scalable interventions, there is one potential LMS analytics variable that has received comparatively little attention: the role of course design.

Part of the problem is how to operationalize something as theoretical, subjective or varied as instructor pedagogy. Indeed, Macfadyen and Dawson [6] attributed variations in "pedagogical intention" as a reason why the LMS could never serve a "one size fits all" dashboard to predict student success across an institution. Similarly, Barber and Sharkey [7] eliminated theoretical student engagement factors such as self-discipline, motivation, locus of control and self-efficacy because they were "not available" (i.e., quantifiable) in the LMS data set, which was their primary object of analysis. Basically, how does one quantify course design that seems qualitatively different from usage log data like logins?

Despite these operational challenges, some of the most frequently cited LMS analytics studies referenced above actually provide a surprisingly uniform characterization of course design that can be roughly broken down into three broad, but distinct categories:

1. User & Content Management (e.g., enrollment, notes, syllabi, handouts, presentations) [1]
2. Interactive tools (e.g., forums, chats, blogs, wikis, announcements)
3. Assessment (e.g., practice quizzes, exams, electronic assignments, grade center use)

If we are willing to accept LMS course design as an aspect of instructor pedagogy – and accept student LMS activity as a proxy for attention, if not engagement – then it may be possible to use one to inform the other. Specifically, patterns of student LMS behavior around tools or functions could retroactively shine light on implemented course design choices that align with the broad, research-based LMS course design types described above.

For example, if students in one course appear to use the online discussion board more than students in another course, could one reasonably assume that instructors of the two courses varied at least in their conceptual value and effective use of this interactive tool? Perhaps this is evident by how instructors differ in their weighting or reward for the discussion board's use in the course's grading scheme, or model and facilitate its use, or simply enable it as a tool in the LMS course's configuration. Admittedly, the challenge is determining how much variance in student LMS course usage is statistically significant or attributable to and indicative of instructor course design. For assessment purposes, though, these three broad LMS course design types (content, interaction and assessment) provide at least a theoretical way to operationalize variability in faculty LMS course design and usage.

While there may be a default institutional LMS course configuration most instructors blindly accept, in trying to explain why one tool or function is used by students more in one course vs. another, it seems odd that we shouldn't be able to consider the pedagogical design choices of the instructor as an environmental factor that may impact student awareness, activity and engagement. True, this may also reflect an instructor's capability or capacity to effectively express his or her pedagogy in the LMS, but to simply ignore the possible impact of course design on student engagement seems un-necessary and disingenuous if we want to use learning analytics to predict and hopefully intervene with struggling students. If students who perform well use the LMS more, do we not want to know what tools, functions and pedagogical practices may facilitate this dynamic?

2. SOLUTION & METHOD

Despite the striking similarity in how several LMS-based analytics studies have categorized LMS course design practices (if

[1] Dawson et al (2008) proposed a 4th type of LMS use called "administration" that roughly equates to course logistics of enrollment, file management, etc. For convenience, I've combined this into the "user & content management" category.

not pedagogical intent), what's needed is a plausible, systematic approach to operationalize these common definitions.

2.1 Weighted Item Count by Design Type

Conveniently, Blackboard used these same research-based definitions of course design for its Analytics for Learn (A4L) product. Specifically, A4L's "course design summary" is a statistical comparison of a Bb course's relative, weighted item count compared to all courses in a department and the institution based on the three major item types found in the LMS analytics literature. Essentially, all items in any Bb course, such as documents or files, discussions or chats, and assignments or quizzes, are grouped into 1) content, 2) interactive tools or 3) assessments. Then, A4L's course design summary uses a simple algorithm to categorize all courses into department and institutional statistical quartiles through the following process:

1. Sum all course items by primary Item Type (e.g., Content, Tools, Assessments).
2. Multiply the group total using a weighting factor (wf): Content (wf = 1), Interaction (wf = 2) and Assessments (wf = 2). [2]
3. Statistically compare each course to all other courses in the department and all other courses across the entire institution.
4. Tag each course with a quartile designation for both the department and institution dimension.

Again, the "course design summary" is already provided in A4L and is really just a way of categorizing how a course is constructed, compared to all courses in the department and across the institution, not necessarily if and how it is actually used by students. To understand and relate student activity to course design, we need to calculate a similar summary of student activity from existing A4L measures.

2.2 Student Activity Summary

Bb Analytics 4 Learn (A4L) contains several student activity measures that include the following:

- Course accesses after initially logging into the LMS;
- Interactions with any part of the course itself, equivalent to "hits" or "clicks";
- Minutes using a particular course (duration tracking ends after 5 minutes of inactivity);
- Submission of assignments, if the instructor uses assignments;
- Discussion forum postings, if the instructor uses discussions.

However, for calculating the companion student activity summary to correlate with A4L's course design summary, I have only used the first three measures (accesses, interactions and minutes) because ALL courses generate this kind of student activity, regardless of design type. Not all instructors use electronic assignments or discussion forums, but short of simply dropping a course, all students generate at least *some* activity that can be measured as logins, clicks or hits and duration.

To calculate the student summary, we must first convert each raw activity measure to a standardized Z-score, which shows how many standard deviations and in which direction a particular raw score is from the mean of that measure in a normal distribution of cases. Because the scale of each activity varies greatly during a semester (e.g., accesses or logins could be under one hundred, interactions or hits could be in hundreds and duration or minutes could be in the thousands), converting these variables to Z-scores allows us to compare and summarize them across measures more efficiently. It also allows us to identify and remove outliers, which for this purpose is defined as scores greater than three (3) standard deviations from the mean. The formula for converting Z-scores is as follows:

$$Z = \frac{X - \mu}{\sigma}$$

The Z-score is equal to X (value of the independent variable) less μ (the value of the class mean for X), divided by σ (the class standard deviation of X).

Accordingly, the steps to analyze and summarize student activity in all courses include the following:

1. Convert accesses, interactions and duration student Bb activity measures to Z-scores.
2. Average the combined student activity scores into a summary measure.
3. Assess the internal consistency of items using a Cronbach alpha test of reliability for each approach (e.g., comparing converted Z-scores).

In addition to student LMS activity and course design measures described above, I used a "threshold" approach to academic performance. Specifically, I used "C or better" final grade in a course and "2.0 or better" term grade point average (GPA) as dependent variables.

3. FINDINGS

3.1 Data

The participants for my study were all first-time, full-time, degree-seeking, undergraduate freshmen or transfer students starting their enrollment in Fall 2013. According to the UMBC Office of Institutional Research and Decision Support (IRADS), this included 2,696 distinct students (1,650 freshmen and 1,046 transfers) or 24.48% of all 11,012 degree-seeking undergraduates.[3] The demographic distribution was as follows:

		(Fresh.%)	(Trans.%)	(Total%)
Gender				
	Male	57	48	54
	Female	43	52	46
	Subtotal	100	100	100

[2] When Blackboard developers were prototyping A4L, I urged them to consider giving "assessments" (e.g., quizzes, surveys, assignments, etc.) a higher weighting of 3, because assessments are more complex for faculty to develop and potentially more impactful on student activity, if not learning. Bb decided not do to this, but does allow A4L's "1-2-2" default weighting to be "customer configurable." We are still evaluating Bb's default weighting, which may be more conservative than my own, but either approach seems reasonable.

[3] http://oir.umbc.edu/files/2013/11/CDS_2013-2014-.pdf

Race			
Asian	24	11	19
Black	11	20	15
Hispanic	5	10	7
White	45	41	43
Other[4]	9	9	9
Unknown	6	10	8
Subtotal	100	100	100

Table 1: Study Sample, 2013-14 FT Freshmen & Transfers

3.2 Grades by Student LMS Activity

Generally, students who performed well academically in courses and a given term overall, showed a higher, statistically significant ($p < .001$) use of Bb compared to peers who did not perform as well. Specifically, using logistic regression to control for other factors such as gender, race, age, Pell eligibility, academic preparation and admit type, students were 1.5 to 2 times more likely to earn a C or better in Fall 2013 and Spring 2014, respectively. Similarly, students were 2.4 to 2.8 times more likely to earn a 2.0 term GPA in Fall 2013 and Spring 2014, respectively.

3.3 Student LMS Activity by Course Design

Generally, students were much more active in Bb courses that used a wider array of Bb functionality. Specifically, after using linear regression, both the institutional course design quartile and instructor use of the grade center were statistically significant ($p < .001$) in terms of freshmen and transfer LMS activity in both semesters. As indicated by the R^2 change, course design and grade center use contributed more than 20% to the overall models, whose adjusted R^2 of .265 and .239 explained 26.5% and 23.9% of the variance in student Bb usage for freshmen and transfers, respectively in Fall 2013. A similar pattern emerged in Spring 2014, with course design and grade center use contributing more than 22% to the overall models' adjusted R^2 of .333 and .278, which explained 33.3% and 27.8% of freshmen and transfer student use of Bb, respectively.

3.4 Student Grades by Course Design

Generally, there was a statistically significant ($p < .001$) relationship for student academic outcomes based on the interaction of course design and student activity in the LMS. However, there was a marked difference in the Expected (B) or odds ratio for both groups of students across both terms, depending on whether I used institutional course design quartiles (ICDQ) or course grade center use as the covariate interaction effect with student Bb activity. For example, the ICDQ * Bb activity interaction effect never produced an odds ratio higher than 1.009, which translates into little more than 1 times the likelihood of earning a C or better final grade (essentially, a 50/50 chance).

By contrast, the odds ratio for the grade center use * Bb activity interaction effect was no less than a 1.571 (for transfers in Spring 2014) and reached a high of 2.455 (for freshmen in Spring 2014). This means that selected subsets of my sample of students had a 1.6 to 2.5 times chance of earning a C or better after controlling for other demographic and academic variables.

Using the same approach for 2.0 or better term GPA, the odds ratio for freshmen under the grade center * Bb activity interaction effect model was 2.610 and 3.504 for Fall 2013 and Spring 2014, respectively. This means freshmen were 2.6 to 3.5 times more likely to earn a 2.0 term GPA in their Bb courses that used the grade center. By contrast, the institutional course design quartile (ICDQ) * Bb interaction effect model remained essentially the same as the C or better findings described above.

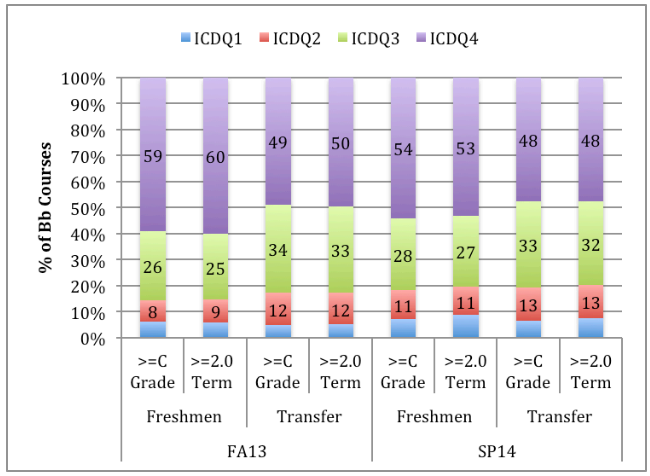

Figure 1: Outcomes by Inst. Course Design Quartile

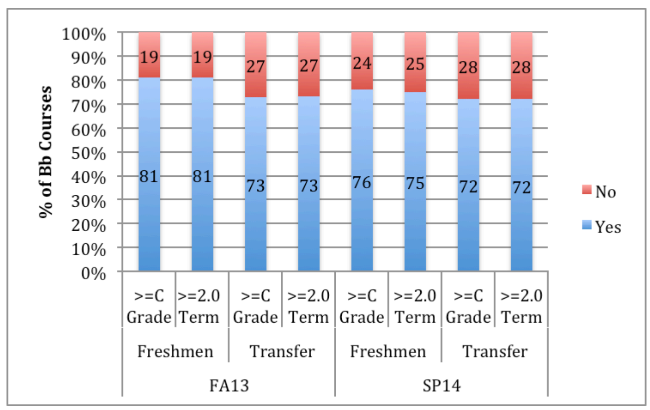

Figure 2: Outcomes by Grade Center Use

4. DISCUSSION

While the correlation between LMS course design and student outcome is compelling, I cannot confirm or reject a hypothesis that it is a causal relationship. I'd want to study these relationships over a longer time, across the entire student population, and even replicate it at other schools. However, is it necessary to establish causality to leverage let alone prove a prediction? Desirable: yes. Necessary: I'm not so sure.

I tend to view LMS use – by faculty and students – as a real-time proxy for their respective attention, engagement and effort in the larger context of teaching and learning. As such, we've developed a simple "Check My Activity" (CMA) feedback tool for students

[4] The "Other" category is my combination of relatively small numbers for "International," "Native American," "Pacific Islander," and "Two or More" UMBC Census Data categories.

allowing them to compare their own LMS activity with peers who earn the same, higher or lower grade for any assignment – provided the instructor uses the grade center. [3] After controlling for other factors (e.g., gender, race, academic prep, Pell eligibility, etc.,) freshmen using the CMA were 1.7 times more likely to earn a C or higher final grade ($p < .001$), but transfers were barely 1 times more likely and the findings were not statistically significant.[5] We also show students how active the LMS course is overall compared to other courses in the discipline, and recently extended this same view to faculty themselves. This way, everyone can decide how to gauge or interpret the importance of their own – or even an entire course's – LMS activity in the context of that exhibited by others.

Additionally, Blackboard has developed a compelling predictive risk model based on this combination of student activity and course design to derive a student "engagement" indicator that is reflected in UMBC's actual full-time freshmen and transfer retention status from Fall 2013 (see figures 3 and 4 below).[6]

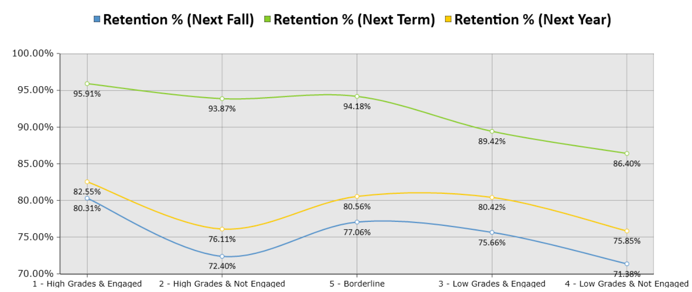

Figure 3: Freshmen Retention by Bb Learn Risk Profile, FA13

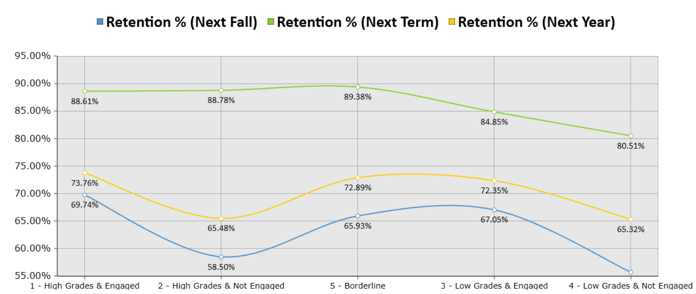

Figure 4: Transfer Retention by Bb Learn Risk Profile, FA13

Notice how less successful but more engaged students (#3) are retained next year at higher rates than more successful but less engaged peers (#2), particularly transfers (figure #4). Moving forward, I can see the Bb integrated model becoming a valuable tool in studying the long-term impact of an LMS on student retention, persistence and graduation. If so, it might also reinforce the value of using the LMS as a real-time indicator of student engagement, not just the passive, one-way delivery of content for which it has typically been used.

[5] Based on my recently defended dissertation available at http://umbc.box.com/johnfritzdissertation.

[6] Larger images and screencast demo available at the following: https://umbc.box.com/fritzpclashortpaperimages

4.1 Course Design as Scalable Intervention

If course design has a relationship with student academic performance, then faculty development could be a necessary first step toward a more scalable form of institutional intervention with at-risk students. In fact, in describing self-directed learning, Ambrose et al [8] suggest that "students must learn to assess the demands of a task, evaluate their own knowledge and skills, plan their approach, monitor their progress, and adjust their strategies as needed" (p. 191). However, instructors also need to be pedagogically ready and secure in their own roles as teachers to desire this kind of empowerment for their students, let alone seek it out by design.

For example, Robertson [9] proposed what is now considered a classic model for how faculty beliefs about teaching influence their evolving pedagogical practice, including the following stages:

- *Egocentrism* – focusing mainly on their role as teachers;
- *Aliocentrism* –focusing mainly on the role of learners; and
- *Systemocentrism* – focusing on the shared role of teachers and learners in a community.

If this evolution of thought and practice occurs at all among teachers, Robertson identifies telltale signs of the transformation. First, as faculty move from one stage to the next, they bring the benefits and biases of the previous stage. Second, they typically change their beliefs and practices only when confronted by the limitations of a current stage, which is brought about by teaching failures. Finally, the desire for certainty, stability and confidence either keeps faculty frozen in a current, status quo framework or drives their progression to the next one in an effort to avoid a potentially paralyzing neutral zone: "a familiar teaching routine that they have deemed inappropriate and with nothing to replace it" (p. 279).

Just as Robertson showed how faculty beliefs about teaching influenced their practice, Steel [10] showed how teaching beliefs influenced their perceptions about what they believe various instructional technologies will allow them to do. For example, using detailed case studies about faculty use of online discussions in an LMS, Steel illustrates the creative tensions between how faculty conceptualize teaching and how they perceive the affordances of web-based technologies like an LMS.

"The velocity of change in the affordances offered by learning technologies presents a significant challenge as does the minimal incentives available to university teachers to use technologies effectively in their teaching practices." (p. 417)

Whether faculty like it or not, when they teach online or use online tools as supplements in their traditional classrooms, they also become webmasters. As such, they need to understand the potential affordances and limitations of web technologies as they attempt to express and implement their pedagogy in course designs. Steel argues that this "reconciliation process" between pedagogical beliefs and rapidly changing technology affordances "needs to be incorporated more fully into informal teacher development approaches as well as formal programs" (p. 417).

To me, faculty who are in Robertson's "neutral zone" between "teaching failures" and "nothing to replace [them]" may be ripe for a course design intervention based on learning analytics, but only if they are aware of peers who they believe have a more effective approach. This is why and how learning analytics may be able to identify, support, promote and evaluate effective practices and practitioners, to serve as a standard by which faculty not only measure themselves, but also point to a way forward, by ideally helping students take responsibility for learning. Yes, technology may help, but per Robertson's and Steel's research, it may not do so unless faculty first *believe* that it can, enough so as to try or look for peers who have done so. Just as students taking responsibility for their learning is the only scalable form of learning, so too must faculty take responsibility for "teaching failures." This includes being open to other pedagogical examples and working hard to master and implement them, which requires a willingness to explore, practice, refine and self-assess.

5. NEXT STEPS

In recent posts, e-Literate bloggers Michael Feldstein and Phil Hill lament the ubiquitous, but essentially boring LMS [11] and even equate it to the minivan of education technology that has long-lasting utility, but not much zip or cache [12]. But if we are willing to go beyond a conventional view of the LMS as more than a content repository or one-way (ego centric?) delivery of knowledge from instructor to student, we might just find that variations in student behavior can shine light on effective course design practices.

Toward this end, we are beginning to look at the LMS as a way to identify effective course design practices and practitioners. While a given semester is underway, we monitor positive outlier courses that appear to generate inordinately high student LMS usage. When the semester is over, we correlate final grades and follow up with instructors whose students may also be performing higher than peers within a department or the institution. To be sure, we conduct these qualitative interviews without necessarily relying on student LMS usage. But taken together, high student LMS usage and grade distribution analysis adds a real-time indicator of student engagement and academic performance that is no longer limited to the end of semester post-mortem.

Finally, as instructional technology support staff, it is not our job to shine light on instructors or course designs that could be better. We've learned instructors learn best from each other, but we can help by using the technology and methodology of learning analytics to identify and reverse engineer effective course design practices we wish all faculty knew about and would emulate. In this way, course redesign could be the most scalable form of analytics-based intervention any institution could pursue.

6. REFERENCES

[1] Campbell, J. (2007). *Utilizing student data within the course management system to determine undergraduate student academic success: An exploratory study*. Retrieved from http://proquest.umi.com/pqdweb?did=1417816411&Fmt=7&clientId=11430&RQT=309&VName=PQD

[2] Dawson, S., McWilliam, E., & Tan, J. P. L. (2008). Teaching smarter: How mining ICT data can inform and improve learning and teaching practice. *Proceedings Ascilite Melbourne 2008*. Retrieved from http://ascilite.org.au/conferences/melbourne08/procs/dawson.pdf

[3] Fritz, J. (2013). *Using analytics at UMBC: Encouraging student responsibility and identifying effective course designs* (Research Bulletin) (p. 11). Louisville, CO: Educause Center for Applied Research. Retrieved from http://www.educause.edu/library/resources/using-analytics-umbc-encouraging-student-responsibility-and-identifying-effective-course-designs

[4] Macfadyen, L. P., & Dawson, S. (2012). Numbers are not enough. Why e-learning analytics failed to inform an institutional strategic plan. *Journal of Educational Technology & Society*, *15*(3), 149–163. Retrieved from http://www.ifets.info/journals/15_3/11.pdf

[5] Whitmer, J. (2012). Logging on to improve achievement: Evaluating the relationship between use of the learning management system, student characteristics, and academic achievement in a hybrid large enrollment undergraduate course. University of California, Davis. Retrieved from http://johnwhitmer.net/dissertation-study/

[6] Macfadyen, L. P., & Dawson, S. (2010). Mining LMS data to develop an "early warning system" for educators: A proof of concept. *Computers & Education*, *54*(2), 588–599. http://doi.org/10.1016/j.compedu.2009.09.008

[7] Barber, R., & Sharkey, M. (2012). Course correction: Using analytics to predict course success. In *Proceedings of the 2nd International Conference on Learning Analytics and Knowledge* (pp. 259–262). New York, NY, USA: ACM. http://doi.org/10.1145/2330601.2330664

[8] Ambrose, S. A., Bridges, M. W., DiPietro, M., Lovett, M. C., & Norman, M. K. (2010). *How learning works: Seven research-based principles for smart teaching*. John Wiley and Sons.

[9] Robertson, D. L. (1999). Professors' perspectives on their teaching: A new construct and developmental model. *Innovative Higher Education*, *23*(4), 271–294. http://doi.org/10.1023/A:1022982907040

[10] Steel, C. (2009). Reconciling university teacher beliefs to create learning designs for LMS environments. *Australasian Journal of Educational Technology*, *25*(3), 399–420. Retrieved from http://www.ascilite.org.au/ajet/ajet25/steel.html

[11] Feldstein, M. (2014, November 10). Dammit, the LMS -. Retrieved from http://mfeldstein.com/dammit-lms/

[12] Hill, P. (2015, May 7). LMS Is The Minivan of Education (and other thoughts from #LILI15) -. Retrieved from http://mfeldstein.com/lms-is-the-minivan-of-education-and-other-thoughts-from-lili15/