Innovating Education with MOOC/FLIP

As I go through rounds of discussion with Princeton University leaders, trustees, alumni, faculty, staff, and students, as well as interested parties from other universities, K-12 schools, and government agencies, it becomes useful to capture some of my experience and thoughts in a brief summary. Much more extensive discussion has been taking place at many levels within Princeton's community, including the faculty committee on online education. *-- Mung Chiang, Jan. 2013*

History and Trends

Higher education is standing at a crossroad today. Massive Open Online Courses (MOOC) and "flipped classrooms" (FLIP) have become two major disruptive forces, creating challenges as substantial as opportunities.

Back in the mid-1990s, with the rise of the web, e-learning in higher education entered into a new phase of development and impact. Online programs, online degrees, and even online universities sprung up and some thrived. In the late 1990s, providers of Learning Management Systems (LMS), such as CollegeNet, Blackboard, and Embark, stabilized as an industry. Early attempts at bringing high quality higher education content to the web also faltered in many cases, partially because the right models, technologies, and consumer behavior were not ready yet in the 1990s.

The focus of this discussion, MOOC, started in 2002 when MIT pioneered Open CourseWare (OCW) despite the many logistic and resource challenges it faced. Many were puzzled by this bold and apparently unnecessary move: "Why would MIT decide to put all the course material online for free? Who would then want to pay the tuition to attend MIT? What does MIT get in return for disseminating knowledge so widely?" It turns out that OCW was an inspiring path-finder that became a major trend. Many top tier universities in the US have joined the movement in the decade following MIT's announcement. Of course, when they talked about course material back then in OCW, they were only referring to lecture notes and homework assignments in PDF files. With the rise of YouTube and Facebook since then, the notion of "course material" would soon change.

In 2007, two other major events unfolded. One is the rise of Khan Academy. How Sal Khan started it on YouTube is now legendary, in part thanks to Bill Gates' introduction to Khan's TED talk that advocated FLIP. Students in K-12 would watch lecture videos before coming to the classroom, where they would get one-on-one interaction with teachers while doing homework. By the count of the number of students who have listened to an instructor's voice, Khan is probably the most heard teacher in human history. Millions of high school students in the US have used Khan Academy. The other event is the creation of iTunes U by Apple. Even though the number of university courses taken on iTunes U has not been as high as Steve Jobs might have hoped for, it marked another step in the MOOC evolution.

Stanford Engineering School has long been recording lecture videos for its (paid) part-time students from the industry in the Bay Area, as a decades-old service and "cash cow." Then Stanford Engineering Everywhere started putting these videos online for free, rather than the "walled garden," paid-only audience. In summer 2011, three computer science courses at Stanford announced that they would put all their lecture videos online for free, together with multiple choice quiz questions that provide instant

feedback at the end of each short video clip, and a social learning environment of discussion forums that were similar to prior platforms like Stack Overflow or Piazza. This whole package, which combined ideas from YouTube-style videos and Facebook-generation social learning, was offered at no charge and to anyone in the world with an Internet connection. The response was overwhelming, with each course enrollment totaling over 100,000 from all over the world.

These numbers rapidly lead to the spin-offs of two for-profit companies backed by venture capital firms in Silicon Valley: Udacity in late 2011 and Coursera in early 2012. MIT started its non-profit MOOC platform called MITx in December 2011, which expanded into the edX consortium that now includes Harvard, Berkeley, University of Texas, Wellesley, and Georgetown University. Within Stanford alone, several other platforms were created in 2012, such as Venture Lab and Class2Go, which consolidated into the aggregation platform of Stanford Online in September 2012. In November 2012, we saw yet another major MOOC consortium of ten universities, including Duke, called Semester Online. Furthermore, other universities such as CMU and Yale have long-standing platforms for online education now in MOOC format. These MOOC platforms are complemented by other non-university-originated sites, including Code Academy, Udemy, Straighter Line, Skill Share, Lore, etc.

Now, in January 2013, at least a dozen MOOC platforms have been established in higher education, most of them within the past two years. Some create new content (e.g., Udacity) and others aggregate existing content (e.g., Coursera). Some are content providers (e.g., Open Yale) while others are also platform providers (e.g., Google Pearson). Some are non-profit (e.g., edX) while others are for-profit (e.g., Coursera). Some develop and use open source platform software (e.g., Class2Go) while others develop proprietary codes (e.g., Coursera). Some issue certificates or even target online diplomas while others allow non-certificate courses. But common across all of them are these features:

- online,
- free,
- open content consumption,
- lecture videos with quizzes and scalable student discussion forums,
- very large numbers of enrolled students (often tens of thousands per course),
- very low completion rates (often 1-3% truly qualify as finishing a course).

My MOOC and FLIP experience

ELE/COS 381: Networks: Friends, Money, and Bytes was a new undergraduate, interdisciplinary course I created in 2011, and became one of the six courses piloted by Princeton on Coursera in 2012: www.network20q.com

Creating this course and carrying out its MOOC/FLIP experiments has been a privilege that I enjoyed tremendously as a Princeton faculty. This opportunity for me to learn so much about learning and teaching was in part supported by Princeton's 250th Anniversary Fund run by Clayton Marsh's office, by Class of '72 Teaching Initiative Fund, by a generous donation by Dr. Bill Ju '78, and by the support from McGraw Center and the Electrical Engineering Department.

This course "Network FMB" adopts the *just-in-time* and *case-study-based* pedagogical approaches to teach the fundamentals of networking, across social, economic, and technology networks. The entire course is structured around 20 practical questions that teenagers today can relate to in their daily lives, questions such as "How does Netflix recommend movies?" "How does Google sell ad spaces?" "What's inside the cloud of iCloud?" and "Why is WiFi slower at hot spots than at home?" Key conceptual themes, such as distributed coordination and network effects, and useful mathematical languages, such as optimization theory and graph theory, are introduced as each of these questions are formulated and answered. Following the just-in-time spirit, only the mathematics that is needed exactly in that lecture is introduced on that spot.

In Fall 2012, over 49,000 students enrolled in this course online, despite its very challenging technical material and focus on concepts (rather than skills), its calculus and linear algebra prerequisites, and the lack of any certificates issued by Princeton. About half of the students were 30 years or older, and one quarter did not have a college degree. Only 30% were from the US, with 35% from Europe and Canada, and 35% from other parts of the world. Half of them were taking this course as their first MOOC, and one third of them came from backgrounds other than science or engineering.

Beyond the standard MOOC features (videos, quizzes, homework, and discussion forums), the TAs and I also experimented with several new methods:

- *Virtual Office Hours* (VOH) provide a scalable broadcast of the TAs and me talking about the course and homework, with a text-based feedback coming in from all the attendants.

- *Grand Challenge Homework* (GCH) provides a pathway for students from any background to demonstrate their unusual technical strength and be connected with top-tier research labs like Bell Labs.

- *Course Wiki* provides a platform for archival knowledge generated by students (through their discoveries and mutual explanations) to be turned into an e-book companion of the course textbook in 2013.

- We also employed several *social media outlets*. For example, our Network20Q Facebook account received 1000 friend requests within a few hours after it was announced, and over 700 Twitter followers showed up on Day 1.

- A variety of interesting *research questions* about online learning efficiency were formulated and tackled based on the empirical data from this course offering.

As the course went online and became available to everyone in the world, including Princeton's own undergrads, it opened the door to redeploy classroom time rather than repeating an identical version of the same lectures. While working on homework together may be a useful model for high school flipped classrooms, for these 30 Princeton undergrads in an elective and demanding course, something else was in order. We used classroom time for three purposes.

- *Question and Answer sessions*: Each class time started with the undergrads asking the instructor targeted questions about the points that they were confused on; twice I could not fully explain the points and had to get back to them after doing more "homework" of my own. I also asked the students many questions, and then based on their answers (anonymously provided through our FLIP website) obtained a very good view of their understandings and misunderstandings, and thus was able to structure "mini-lectures" accordingly targeting the confusing points. A

basic version of the software to enhance collaboration and discussion was also created, including the feature for anonymous response to questions in light of the "silent classroom syndrome," where the vast majority of the students never raise hand to speak up.

- *Guest lecturers from industry*: 15 prominent leaders from the industry (and some from government and academia) were invited to provide guest lectures and answer questions from the students. A panel of venture capitalists was also convened in one class for the entrepreneurship-oriented students.

- *Real-time demonstrations and experiments*: From "WiFi packet capture" to "greedy social network forwarding," we carried out many rounds of demos of social and technological network ideas. We also spent one class on a round-robin discussion, where each student thought like an AT&T CTO and proposed solutions to the mobile data demand growth.

While experimenting with what FLIP should look like for an engineering elective course, there were also somewhat unexpected lessons for me. For example, during one class, I started to feel that the percentage of students getting the wrong answer was unusually high. So I asked a binary question for anonymous responses: "have you actually watched the video lecture on YouTube or Coursera before coming to class today?" It turned out that 90% of them had not. It was the day before most of their midterm exams. The spine of synchronized learning weakened. Do we have something new to strengthen it or replace it?

Opportunities and Challenges: 10 Questions

The variety of MOOC platforms, and the inevitable implications of FLIP, has presented remarkable opportunities to redraw the boundaries of education. There is also no shortage of challenges, as substantial as the opportunities we face. The list of 10 questions below highlights both a rosy picture that is irresistibly attractive and many obstacles and uncertainties along the path.

1. *Who can access education and opportunities globally?* For centuries, higher education institutions' mission has started with a set of students who *dare* to apply. And that set grew dramatically several times over the past century. Can these institutions revisit the mission now, to consider it as the search for talents around the world who *should* apply? The impact of this question spans K-12, higher education, enterprise and vocational training, and all the way to lifelong learning. With the advent of high quality education being made widely accessible through technology, we may be standing at the cusp of another major wave of democratization of education.

2. *Can American colleges become more affordable?* On what grounds can ever-increasing tuition, 3% above inflation rate for decades, be charged? Parents and students will demand answers to this question, starting perhaps from lower-tier private universities and those public universities that charge hefty out-of-state tuition, but certainly they will keep asking the question to all the educational institutions. More broadly, what is a university offering to its students through its education on campus anyway (other than branding and social networks that are also increasingly independent of the physical campus)? Princeton's residential colleges and similar systems at other universities may shine brighter than ever before.

3. *What is the role of the classroom and of faculty?* Should class be reserved only for two-way interactions, like what many law schools, business schools, and humanity departments have

been doing, as one-way, open-loop, behind-auditorium-door lecturing moves onto YouTube and MOOC? Will online education's scale incentivize teachers to pay even more attention to every word they utter and record for the world to hear? And what constitutes "teaching outside of a university" now?

4. *Is there a sustainable revenue model for MOOC?* Platforms like Coursera, Udacity, and Code Academy are profit-seeking companies financed by venture capitalists seeking returns on their investment, while those like EdX, Class2Go, Semester Online, Open Yale are non-profits financed by university resources. Either way, financial sustainability of some degree is needed. Recent articles in the WSJ and the NYT have reported the universal consensus today: there is yet no business model to generate non-trivial amounts of revenue for profit-seeking MOOC platforms.

5. *How can MOOCs assess and authenticate students?* There are at least four variants of student assessment: (a) self-grading, (b) machine-grading, (c) peer-grading, and (d) professional-grading. It is difficult to use the first three types of grading for certain subjects, yet the last type is not scalable, contradicting the fundamental premise of "self-scaling" in MOOC. Regarding student authentication, there are already profitable companies selling Coursera certificates for a fee, and they have ``solutions" to defeat the recently announced biometric authentication tests. Maybe education can be free, but must authentication be done with test centers and ground staff (with some fee)? Furthermore, what kind of certificate should be given? There are at least four different types of certificates, each demanding different confidence levels of authentication: (a) a diploma, (b) letter grades, (c) pass/fail, or (d) a personal letter from the instructor, bearing no institutional logo, congratulating the student on completion of basic requirements.

6. *How can we rate, recommend, and assign credentials to courses?* This is the mirror image of the student-assessment challenge above. How do students, parents, and employers know what a course has taught? What does an A- mean for one course relative to a B+ from another? Is there consistency of quality control from one semester to another? Which courses should students take after completing the current one?

7. *What is the representation and ontology of knowledge online?* In order to have a truly meaningful and somewhat universal credentialing of online courses, such an ontology will be essential. On this fundamental challenge, we are still light years away from turning "open online courses" to structured "open online education," from MOOC to *MOOE*.

8. *How can MOOCs scale up, scale down, and engage students?* Are students actually paying enough attention in asynchronized learning? How big can *effective* social learning become, if there is an inevitable tradeoff between the enrollment size and the effectiveness of individual learning? Can we create personalized instruction, perhaps precisely because of the massive scale, just like how Netflix uses collaborative filtering (leveraging similarities of movie tastes among people) to better recommend movies to each individual? Can textbooks of the future learn each reader's learning style from data collected thus far, and adapt both the content and layout of the next chapters? (see http://scenic.princeton.edu/latestNews/book_021413.html)

9. *How might we make it easier for faculty (or any articulate expert) to start a course?* How much time does it really take? How many resources does it take to produce and maintain a course? Sal Khan did not recruit a camera crew to start Khan Academy; his recording setup costs $250 plus a tablet. A mobile recording station costs in the order of $5000, but takes only a desk space and

can be shared by many teachers. For a university like Princeton, time should be the resource bottleneck of MOOC development rather than funding. Many faculty members are hesitatant to commit a whole summer to create a semester-long MOOC, nor do most of the MOOC students have the luxury to engage in semester-long learning. Perhaps the semester should not be the temporal granularity for the new learning environment, and we can avoid fitting the content format from the old media into the new media. (see www.3nightsdone.org)

10. *How might we make it easier for students to finish a course?* Completion rates for MOOCs today are typically on the order of 3% (and more rigorous assessment methods in the future may reveal even this number as an over-estimate). In hindsight, that is hardly surprising. Why would, say, a 40-year old, who has a spouse, two kids, one dog, and two mortgages, put in 20 hours a week on difficult homework about differential equations and then feel more miserable in life? Signing up for a course takes 30 seconds and provides an emotional boost. In contrast, completing 12 weeks of a rigorous college course suffers from the "new year resolution syndrome." People with limited resources and background suffer most. More than just boosting the impressive enrollment numbers further, how can we significantly increase the completion rate (while recognizing that a completion rate much above 20% may be neither feasible nor desirable if rigorous learning is to be carried out in today's MOOC demographics)? For example, certificate in style (d) in item 5 above may provide a *sense of urgency* in completing a course without the burden of student authentication or brand dilution.

Many of the above challenges stem from the fact that a MOOC is intrinsically *asynchronous* and *heterogeneous*. There is no common timetable, no common daily schedule, no common location, no common education goals, and no common background among the diverse demographics in MOOC. The common denominator of higher education dissolves in MOOC.

Where is the Limit? The Law of G*

This asynchronous and heterogeneous mode of learning has lead to what I would describe as a "law," not of nature by laws of physics, but a shorthand for some persistent observations. Given one course offering (i.e., one course over one semester, with one teacher and a few assistants), let E denote the enrollment size of a class, Y denote the yield rate (percentage of students who actually understand the material well enough to deserve an A, B, or C grade), and G denote the Graduation size. Obviously,

G = E x Y

Now, how big can G be, if we innovate the mode of delivering the content? Let's call this upper limit of G the best Graduation number achievable, G*.

The Law of Conservation: G* remains a constant (in order-of-magnitude) for a given type of course.

The order of G* may change for different types of courses, but remains the same for a given course type no matter how we scale the teaching and learning, due to the intrinsic tradeoff between E and Y. For example, for an introductory computer science course, G*=1000. For an upper-class engineering elective, G*=100. In a traditional classroom, E may be 120, Y is about 80%. In a walled-garden continuing education system, E maybe 1000, Y is about 10%. In a MOOC, E maybe 20,000, Y is about 1%.

What the Law of Conservation on G* says is that learning is a human activity with a hard-to-break limit on its "automat-ability" to massive scale. This law has significant implications to MOOC students' welfare, MOOC providers' business models, and colleges' strategic plans. The good news is that even

doubling or tripling G* (within the same constant order-of-magnitude) is still helpful, let alone the potential that those who do not “graduate” may still see more benefit than harm by enrolling in any modality of learning.

At the Crossroads

We are blessed to be “in the arena” as the tsunami of change arrives on the shore of education. For teachers and students alike, this is the most exciting time. This is also the scariest time. As we keep reading media reports of the growth charts of MOOC enrollment numbers, the balance between enthusiastic creativity and cautionary prudence has become even more delicate.

We did not arrive at the current mode of teaching and learning in just one year. It took decades, if not centuries, for it to stabilize. And it will take at least many years beyond 2013 for something new to become something mature in how we teach and how we learn. For certain fields in social sciences and the humanities, it will take even more time and creativity as the automation and scaling-up of education introduce special challenges to individualized, face-to-face mode of learning.

There will be more dead-ends and false starts along the path of this experiment than there is patience and goodwill across the diverse sectors in society. Different participants in this ongoing experiment have vastly different expectations, some of which have now become so great that disappointment is almost inevitable, leaving bad taste in the mouth in future generations (similar to those when we talk about the dot com bubble and the tulip mania). And yet a “worst-case mentality” is hardly a reason strong enough to hold us back and watch the game from the bench.

So we have to ask ourselves: “where is the bar that we are targeting?” If the bar is to completely replace traditional face-to-face learning and teaching, we are bound to miss the goal. If the bar is to give an option of learning to the youths next to the option of playing Angry Bird, we are destined to succeed. There is no consensus today on where that bar will be or exactly what can be achieved, for that future lies, to a large degree, in our hands today.