

Demographic Differences in How Students Navigate Through MOOCs

Philip J. Guo

MIT CSAIL / University of Rochester
pg@cs.rochester.edu

Katharina Reinecke

University of Michigan
reinecke@umich.edu

ABSTRACT

The current generation of Massive Open Online Courses (MOOCs) attract a diverse student audience from all age groups and over 196 countries around the world. Researchers, educators, and the general public have recently become interested in how the learning experience in MOOCs differs from that in traditional courses. A major component of the learning experience is how students navigate through course content.

This paper presents an empirical study of how students navigate through MOOCs, and is, to our knowledge, the first to investigate how navigation strategies differ by demographics such as age and country of origin. We performed data analysis on the activities of 140,546 students in four edX MOOCs and found that certificate earners skip on average 22% of the course content, that they frequently employ non-linear navigation by jumping backward to earlier lecture sequences, and that older students and those from countries with lower student-teacher ratios are more comprehensive and non-linear when navigating through the course.

From these findings, we suggest design recommendations such as for MOOC platforms to develop more detailed forms of certification that incentivize students to deeply engage with the content rather than just doing the minimum necessary to earn a passing grade. Finally, to enable other researchers to reproduce and build upon our findings, we have made our data set and analysis scripts publicly available.

Author Keywords

Massive Open Online Course; MOOC; non-linear learning; navigation strategy

ACM Classification Keywords

H.5.1. Information Interfaces and Presentation (e.g. HCI): Multimedia Information Systems

INTRODUCTION

Massive Open Online Courses (MOOCs) are beginning to globalize education by enabling participation from students of all age groups and nationalities. Previous studies have shown that MOOC students come from over 196 countries, speak a wide variety of languages, have a large range of age

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
L@S 2014, March 4–5, 2014, Atlanta, Georgia, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2669-8/14/03...\$15.00.

<http://dx.doi.org/10.1145/2556325.2566247>

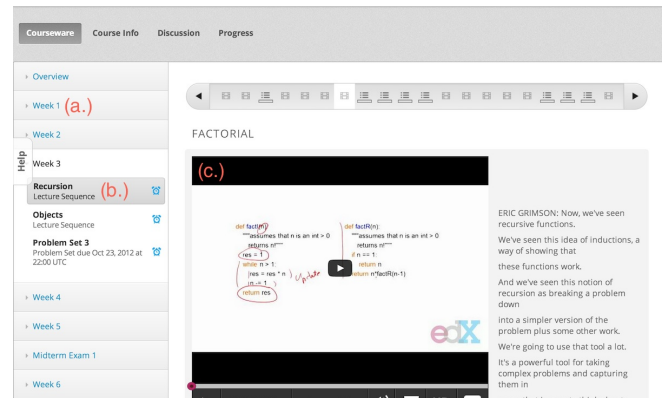


Figure 1. An online course on the edX platform has a hierarchical structure containing a series of weeks (a.), where each week contains several web pages (b.), and each page is a learning sequence made up of video lectures and/or assessment problems (c.).

and prior education, and highly heterogeneous motivations for enrolling in online courses [3].

This diversity challenges the homogeneous learning environments offered by the current generation of popular MOOC platforms such as Coursera, edX, and Udacity (often called xMOOCs [10]). In particular, critics have warned that MOOCs cannot cater to individual and cultural variabilities in learning style, but instead require every student to be a “self-directed autodidact” [13] in order to succeed.

This sort of self-guided learning that current MOOCs require potentially make them ill-suited for students who are less efficient at defining a learning path by themselves [18]. Researchers have found that some learners prefer either more linear or non-linear navigation [5, 17, 21], and some perform best under highly structured guidance [9, 15]. However, current MOOCs are often organized in linear ways with weekly video lectures and assessments resembling the conventional university classroom.

If students have varying needs and preferences for navigating learning content, we should be able to observe differences in how they interact with MOOCs. This paper contributes a study of students’ strategies for navigating through the learning content in MOOCs, and is, to our knowledge, the first to investigate demographic differences. In particular, we hypothesized that there should be notable differences in how students navigate through the learning content depending on demographics such as age and country of origin.

Certificate earners view only 78% of learning sequences, on average; they completely skip 22% of course content.

Certificate earners engage in non-linear navigation behavior, often jumping backward to revisit earlier lectures.

Navigation backjumps from assessments to lectures are more common than lecture-to-lecture backjumps.

Older students and those from countries with lower student-teacher ratios (e.g., the US and European countries) visit and repeat more lecture sequences, which indicates more non-linear navigation and learning strategies.

Younger students and those from countries with higher student-teacher ratios (e.g., India, Kenya) visit and repeat fewer sequences, which indicates more linear navigation.

However, the effect of age is stronger than that of country; older students from countries with higher student-teacher ratios behave more like their similarly-aged counterparts in countries with lower student-teacher ratios.

Table 1. Summary of the main findings that we present in this paper.

To evaluate these hypotheses, we analyzed student interaction log data from four different MOOCs provided by edX, comprising 140,546 students from 196 countries.

Our results show that there are significant differences in how students approach the learning content. Older certificate-earning students cover more learning materials and navigate the learning content in a more non-linear way than younger students. We also found differences between countries, with students from the U.S. and many Western European countries covering and repeating more learning sequences — indicating non-linear navigation — than students from places such as India or Kenya. Finally, our results showed that independent of demographic background, the navigation behaviors of certificate-earning students are frequently motivated by opportunistically working backward from assessment questions.

This paper makes the following contributions:

- We present the first analysis of how demographics affect students' navigation behaviors in MOOCs. Table 1 summarizes our main findings.
- Based on the findings, we provide design recommendations for MOOCs to better support students in understanding learning goals and achieving higher content coverage.
- We make our anonymized dataset and all analysis scripts public (available at <http://www.pgbovine.net/edX/>). To our knowledge, this is one of the first open data sets on MOOC demographics and student activities.

We begin by introducing MOOC platforms and reviewing the literature on different learning styles that impact students' navigation strategies. The main part of this paper includes our study methodology, findings, as well as a discussion of

the results from which we infer design implications. We close with study limitations and opportunities for future work.

BACKGROUND ON MASSIVE OPEN ONLINE COURSES

The research community is currently studying two main types of MOOCs: cMOOCs take a connectivist approach to teaching, emphasizing peer-based social learning and knowledge generation [6, 19]. xMOOCs aim to scale up traditional lecture-based courses by offering online video presentations and exams [3, 7, 10]. This paper focuses on xMOOCs.

Organizations such as Coursera, edX, and Udacity offer free xMOOCs, which comprise a linear arrangement of short video lectures and assessments that are either graded automatically or via peer grading. While xMOOCs offer discussion forums, their structure is mostly teacher-centered, with the instructor creating most of the instructional content [10]. However, xMOOCs differ from in-person courses in important ways: Students can learn at their own pace, repeat or skip lessons, and need to drive the learning process more independently than in offline learning environments.

On the edX platform, each course is divided into a series of weeks (usually 8 to 16 weeks), which are listed vertically on the left pane of the web interface (see Figure 1). Each week contains several web pages called *learning sequences*. Each learning sequence is a web page that contains either a *lecture* or a graded *assessment*. A lecture sequence is a series of instructional videos with optional interstitial quiz problems in between videos. An assessment sequence is either a weekly problem set or a midterm or final exam. Although materials are released weekly and presented in a linear and chronological structure, students are free to navigate back to already published content, even when taking an assessment.

In the remainder of this paper, we use the terms “learning sequence” or “sequence” to mean a single course web page containing learning materials.

DIFFERENCES IN ONLINE LEARNING STRATEGIES

xMOOCs have been criticized for simply replicating traditional lecture-based teaching [2] and not catering to different learning strategies, such as a student's preference for either a more or less linear structure of instructional content [20].

Differences in students' approaches to learning are often described with the help of Witkin's distinction between field-independent and field-dependent learners [24]. Field-independent learners predominantly approach the learning content in an analytic manner, focusing first on details before abstracting from a specific problem. In contrast, field-dependent learners first reason about context before focusing on details [24].

Both learning styles have been investigated in the context of online learning. Field-independent learners are believed to be fairly confident in defining their own learning paths and navigating in non-linear learning environments [9]. They are also often referred to as “explorers”, indicating that they navigate more freely without necessarily following the path suggested by content creators [17]. Field-dependent learners, in contrast, prefer following an externally defined learning path, as

Course	Subject	University	All Students		Certificate Earners	
			total	with demographics	total	with demographics
6.00x	Intro. CS & Programming	MIT	65,475	50,581	5,758	4,155
PH207x	Statistics for Public Health	Harvard	31,851	28,363	4,915	4,528
CS188.1x	Artificial Intelligence	Berkeley	24,517	17,066	1,900	1,232
3.091x	Solid State Chemistry	MIT	18,703	14,152	2,072	1,575
Total			140,546	110,162	14,645	11,490

Table 2. The numbers of students and certificate earners for four edX courses in Fall 2012. The “with demographics” columns show the numbers of students in each group who filled out all demographic information (birth year, education level, and gender), which comprises ~78% of students.

imposed by a teacher or the online learning environment [9]. Often described as “observers” [17], their preference for linear learning has also been found to result in more disorientation problems when presented with non-linear teaching materials in an online context [4]. This trait could be a reason why field-dependent learners spend more time on navigation and used sequential steps (back or forward buttons) less frequently than field-independent learners [16].

Which learning style students adopt is partly influenced by the form of education they received as children, for example in kindergarten or primary school [12]. The predominant form of educational exposure in a society depends on which instructional style that society values most, but also on its financial resources, which directly determine class sizes and therefore, the *student-teacher ratio* [23]. Researchers have found that, in countries with a higher student-teacher ratio, students are more accustomed to a teacher-centered education and behave more like “observers” than “explorers” [11].

However, researchers have also pointed out that learning styles are dynamic, so that learners might adopt new strategies when required [12] or as they grow older [8].

Given these previous findings, we hypothesize that the diversity among students participating in MOOCs will result in notable differences in their navigation strategies.

METHODOLOGY

To understand students’ navigation strategies in MOOCs, and whether their approaches might depend on demographics, we performed a quantitative data analysis on the activities of 140,546 students in four online courses on the edX platform.

Data Set

Table 2 shows an overview of the data we extracted from four courses in the first edX batch offered in Fall 2012. We selected courses from all three edX affiliates at the time (MIT, Harvard, and UC Berkeley). To maximize diversity in subject matter and student population, we selected an introductory computer science course (6.00x), an advanced computer science course (CS188.1x), statistics for public health (PH207x), and solid state chemistry (3.091x). EdX launched additional courses in Spring 2013, including some humanities ones, but that data was incomplete when we began this study.

We have made an anonymized version of our data set and analysis scripts publicly available at <http://www.pgbovine.net/edx/> so that anyone can reproduce and build upon this paper’s findings.

Analysis Variables

We created analysis variables for each student based on their demographics and interactions with the courseware.

Demographics

When students first register for a free edX account, they can enter optional information such as their birth year, highest level of education completed (elementary school, junior high, high school, bachelor’s, master’s, Ph.D), and gender. Across all four courses, 78% of students filled out all demographic information. For analyses that require demographics, we excluded students who did not fill out the relevant fields. We derived additional demographic variables for each student:

- **Age** during Fall 2012, estimated from birth year. We filter out ages that are less than 10 or greater than 80 years old.
- **Years of education** – A numerical estimate based on the student’s highest level of education completed. For example, “high school” translates into 13 years of education.
- **Country of origin**, determined by looking up student IP addresses in the MaxMind GeoIP database [1]. If a student accessed the course website from IP addresses in multiple countries, then we use the country with the most accesses.
- **Student-teacher ratio** in the country of origin – For each country, we obtained the most recent primary school student-teacher ratio (total number of students divided by number of teachers) from the UNESCO Institute for Statistics [23]. This is one widely-used indicator of educational quality and level of individualized student attention. However, note that this variable is usually correlated with economic indicators such as per-capita GDP and median household income, so observed effects might not be directly due to pedagogical quality.

Motivation & Intent

For each student, we extract the following variables that indicate their motivation and intent for enrolling in MOOCs:

- **Certificate earned** – whether a student earned a *certificate of completion* or not. EdX gives out certificates to all students who earned above a passing grade – usually around 60% – set by the course instructor. Students who earn certificates all intended to engage with lectures, problem sets, and exams, and persisted through most of the course. On the other hand, students who did not earn a certificate *and* did not attempt the assessment problems might have been casual bystanders, auditors, or early dropouts [14]. Thus, certificate earned serves as a control variable for student intent. Many of our analyses only consider students who

earned certificates (~ 10% of all students), to focus on those who intended to engage seriously with the course.

- **Grade** – Students earn a grade between 0% and 100% depending on their performance on problem sets and exams. Grades indicate student knowledge and also engagement. Once a student passes the threshold for earning a certificate, higher grades indicate more self-motivation to learn, while lower grades might indicate that a student just wants to do the bare minimum to earn a certificate.
- **Coverage** – The fraction of total learning sequences (lectures, problem sets, and exams) that the student visited.
- **Discussion forum events** – The number of times a student posted to the discussion forum, divided by that student’s total number of courseware access events, which controls for variability in student activity levels. Forum participation has been found to be an indicator of social engagement in MOOCs [7].

Navigation

We quantify the following kinds of non-linear navigation through the course materials:

- **Backjumps** – The number of times that this student navigated backwards from a learning sequence to another one released earlier in the term (e.g., from Lecture 6 to Lecture 4), divided by the total number of sequences visited by this student.
- **Textbook events** – The number of times that this student accessed the digital textbook associated with the course, divided by the student’s total number of courseware access events. Since the textbook still resides within the edX website but lies outside the main flow of a course, we count textbook events as an instance of non-linear navigation.

Analyses

To assess whether the aforementioned analysis variables had statistically significant and independent effects on navigation-related metrics, we conducted *multiple linear regression* analyses and report their ANOVA F statistics, *p*-values, and, when applicable, the unstandardized *b* coefficients for each independent variable in the regression.

For this paper, we do not analyze fine-grained navigation within a learning sequence. We also do not use time as a feature since it is hard to determine exactly how much time a student was actively interacting with particular courseware resources solely from analyzing the server logs we were given.

FINDINGS

Here we present our findings, starting with an overview of student demographics, motivation, intent, and how each affect student navigation through the edX course materials.

Overview of Student Population

The mean student age across all four courses in our data set was 28 years (sd=9.4). Most students (77%) were between 20 and 40 years old, with 13% under 20 and 10% over 40.

The most common highest education level for students participating in these four courses was a bachelor’s degree (38%),




Figure 2. Distributions of self-reported education levels for all students (top) and certificate earners (bottom) in all four courses. * represents an elementary or junior high school graduate. Certificate earners tend to have more years of education than the general student population.

	6.00x	PH207x	CS188.1x	3.091x
	123 countries	133 countries	103 countries	106 countries
	U.S. (22%)	India (17%)	U.S. (19%)	U.S. (18%)
	India (9%)	U.S. (16%)	India (9%)	India (11%)
	Russia (7%)	Spain (8%)	Russia (8%)	Spain (11%)
	Spain (6%)	U.K. (5%)	Spain (8%)	U.K. (6%)
	U.K. (6%)	Germany (3%)	U.K. (6%)	Russia (5%)
	other (50%)	other (51%)	other (50%)	other (49%)

Table 3. The top five countries with the most certificate-earning students in each course, and the percentage breakdown of students from each country. The top of each column shows the total number of countries with certificate earners for each course. Note that approximately half of all certificate earners in each course came from the top five countries.

followed by a high school diploma (28%). However, students who earned a certificate most commonly held a master’s degree (37%), followed by a bachelor’s degree (32%). Figure 2 provides an overview of the distributions of highest education levels, showing that students with higher education levels were more likely to earn certificates.

Most students in these four courses were men, with the largest gender disparity in the two computer science courses: 86% male in CS188.1x (artificial intelligence), 83% in 6.00x (introductory computer science), 70% in 3.091x (chemistry) and 56% in PH207x (statistics for public health).

Students from all 196 countries (plus 10 additional dependent territories such as Guernsey) participated in the four courses. Students from 157 countries ended up getting certificates. Table 3 shows their breakdown by country, with the U.S. and India most highly represented. In contrast, major East Asian countries such as China, Japan, and Korea are notably underrepresented, making up only 0.1%, 0.08%, and 0.07% of the total student population (i.e., including those who did not earn certificates), respectively.

We also found that age distributions vary by country. In particular, students from countries with lower student-teacher ratios seem to participate in MOOCs later in life than students from countries with higher student-teacher ratios (i.e., larger class sizes). Spearman’s correlation coefficients between a student’s age and the student-teacher ratio of their home country were $r = -.25$ for 3.091x, $-.24$ for CS188.1x,

	6.00x, $R^2 = .14$		PH207x, $R^2 = .13$		CS188.1x, $R^2 = .15$		3.091x, $R^2 = .18$	
	coefficient	p -value	coefficient	p -value	coefficient	p -value	coefficient	p -value
Age	.003	< .001	.004	< .001	.001	< .001	.01	< .001
Student-teacher ratio of country	-.002	< .001	-.004	< .001	-.001	< .001	-.004	< .001
Gender (is male)	.03	< .001	.00	.14	.04	< .001	.03	.02
Grade	.31	< .001	.57	< .001	.18	< .001	.36	< .001
Constant	.53	< .001	.19	< .001	.61	< .001	.21	< .001

Table 5. Multiple linear regression analyses of student demographics and grades versus coverage for certificate earners. We excluded years of education from the analysis due to its strong collinearity with age. Across all four courses, a student’s age and grade are positively correlated with coverage, while the student-teacher ratio of the student’s country is negatively correlated.

	6.00x			PH207x		
	df	F	p	df	F	p
Age	1	130	< .001	1	238	< .001
Country of origin	114	4.4	< .001	128	4.5	< .001
Gender	1	23	< .001	1	1.1	.3
Years of education	1	1.3	.25	1	18	< .001

	CS188.1x			3.091x		
	df	F	p	df	F	p
Age	1	35	< .001	1	160	< .001
Country of origin	87	2.3	< .001	97	3.7	< .001
Gender	1	30	< .001	1	1.8	.18
Years of education	1	.5	.49	1	.5	.50

Table 4. Multiple linear regression analyses of student demographics versus coverage for certificate earners. We report ANOVA F statistics, degrees of freedom (df), and p -values. Age and country have statistically significant effects on coverage across all four courses, while gender and education years do not.

-.22 for 6.00x, and -.08 for PH207x, all $p < .001$. For instance, the U.S. has a low student-teacher ratio of 14, and its mean student age is 33 years old ($sd=12$); in contrast, India has a high student-teacher ratio of 40, and its mean student age is 27 ($sd=9$). One possible interpretation is that people from countries with high student-teacher ratios are more likely using MOOCs to supplement their regular education, whereas those from low student-teacher ratio countries – U.S., Canada, Western Europe – are more likely to be adult lifelong learners.

Motivation & Intent

While demographics – most notably age and country – point to potential differences in students’ motivations for participating in MOOCs, we also looked at the amount of learning content that certificate earners covered, what grades they received, and how much they participated in discussion forums.

Coverage

To understand students’ strategies for earning certificates, we first analyzed how much of the course materials (i.e., learning sequences) certificate earners covered. Certificate-earning students viewed, on average, 67% of the learning sequences in 3.091x, 77% in PH207x, 81% in CS188.1x, and 86% in 6.00x. Thus, students ignore, on average, 22% of the materials in those courses, yet still earn certificates. While we are




Figure 3. Age versus mean coverage for certificate-earning students in all four courses. Error bars represent the standard error of each age group’s mean. The oldest group (over 50 years old) achieves, on average, 10% more coverage than the youngest group (under 20 years old).

not able to determine how deeply they engaged with the content that they accessed, coverage measurements at least show that a non-trivial amount of content gets completely skipped.

To evaluate the effects of demographics on coverage, we conducted multiple linear regressions with the demographic factors age, country, gender, and years of education as the independent variables and coverage as the dependent variable. We found that age and country of origin have significant effects on the fraction of sequences that certificate earners cover in all four courses (see Table 4 for a summary of the F statistics).

Age is positively correlated with coverage, even when accounting for other demographic variables (regression coefficients $b = .001 - .01$, $p < .001$ for all four courses¹). Figure 3 visualizes how older certificate earners cover more of the learning content, with a 10% difference in coverage between the under-20 and the over-50-year-old groups. (For this figure and related figures, we binned ages into groups because the data was too sparse for certain ages.)

Table 4 also shows that a student’s country of origin has a statistically significant effect on their coverage. As an example, consider the top two most represented countries: Certificate

¹The unstandardized model coefficients b reported in this paper are small because dependent variables (e.g., coverage) are often between 0 and 1, so independent variables are scaled down by b .

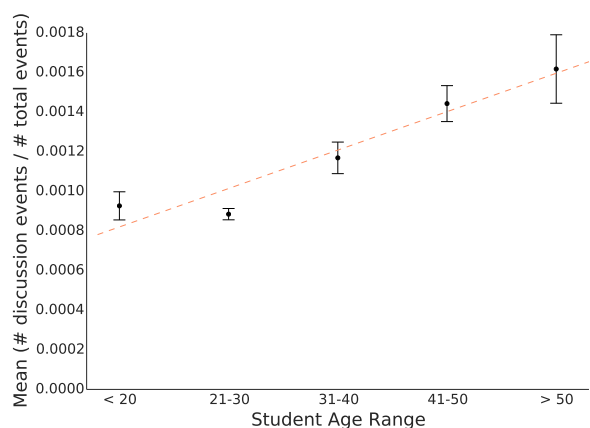


Figure 4. The mean number of normalized discussion forum posting events for all certificate-earning students in each age group. Error bars represent the standard error of each age group’s mean. In general, older students post relatively more on course discussion forums.

earners from the U.S., on average, cover significantly more sequences (83%) than Indian certificate earners do (71%) (independent t-test, $T_{(2829)} = 20.2, p < .001$).

We followed up this finding by investigating the relationship between coverage and the student-teacher ratio of each country. We conducted a similar multiple linear regression analysis with coverage once again as the dependent variable and demographic factors and grades as independent variables. The analysis output in Table 5 shows that certificate-earning students from countries with higher student-teacher ratios usually visit fewer learning sequences.

Grades

Table 5 shows that grades are positively correlated with coverage. In all four courses, students who viewed more materials were more likely to achieve higher grades, which makes sense because they had more opportunities to learn and to get assessed (i.e., higher “time-on-task”). To investigate whether demographics also affect grades, even controlling for coverage, we conducted multiple regressions of demographics and coverage on grades. When the regression used data from all students, it showed that age and student-teacher ratio have significant effects on grades in several courses ($F_{(1)} = 47 - 404$ for age in CS188.1x, PH207x, and 6.00x, $F_{(1)} = 973 - 2402$ for student-teacher ratio in CS188.1x and 6.00x, all $p < .001$). However, the effects of age and student-teacher ratio are non-significant in three of the four courses when considering only certificate earners in the analysis; the only exception is CS188.1x, where $F_{(1)} = 11$ for age, $F_{(1)} = 36$ for student-teacher ratio, with $p < .001$. Thus, once students pass the certificate-earning bar, then grades no longer differ as much across demographics.

Discussion forum posting

An alternative measure of student motivation is how actively they participate in the course discussion forum. We conducted a multiple linear regression with normalized forum events as the dependent variable and demographics as independent variables. The analysis shows that age and




Figure 5. The frequency of each kind of backjump for certificate-earning students. The most common kind was jumping from an assessment (i.e., problem set or exam) back to an earlier lecture sequence.

level of education had significant effects in three of the four courses, with the overall analysis $F_{(2,1572)} = 4.4$ for 3.091x, $F_{(2,4525)} = 27.4$ for PH207x, $F_{(2,4152)} = 24.6$ for 6.00x, all $p < .001$.

Figure 4 shows that older students participate more in forums. And even controlling for age (using the aforementioned regression), students with Ph.D. degrees participated up to 32% less in forums than non-Ph.D. holders. Finally, we expected students from more teacher-centered educational systems (i.e., higher student-teacher ratio) to be less likely to participate socially in MOOCs, but that turned out not to be the case. Student-teacher ratio did not have any significant effects on forum participation in our regressions, and Spearman’s correlation was nearly zero ($r < .05$ for all courses, with $p > .001$).

Course Navigation Strategy

Given that all certificate earners had a common intent – to earn a passing grade – we wanted to understand how they went about doing so, and whether navigation strategies differed by demographic. Specifically, we analyzed one salient form of non-linear course navigation: a *backjump* from one learning sequence to another sequence released earlier in the term. We focused on backjumps because the ability to go “back in time” to view prior lectures or to re-try old assessments is one key differentiator of MOOCs over traditional residential courses with in-person lectures and exams.

Certificate earners in our four courses performed an average of 1.04 backjumps for every learning sequence they visited (sd=0.62). This behavior indicates that students apply non-linear navigation strategies on their way to earning certificates. In contrast, students who did not earn certificates performed only 0.3 backjumps per visited sequence (sd=0.4). Thus, certificate earners repeat visiting prior sequences three times as often, presumably to review older content.

Kinds of backjumps

Each backjump can start and end at either an assessment (i.e., problem set or exam) or a lecture. Figure 5 visualizes a per-

	6.00x			PH207x		
	df	F	<i>p</i>	df	F	<i>p</i>
Age	1	26	< .001	1	62	< .001
Country of origin	114	1.5	< .001	128	2.7	< .001
Gender	1	56	< .001	1	45	< .001
Years of education	1	.42	.52	1	4.7	.03

	CS188.1x			3.091x		
	df	F	<i>p</i>	df	F	<i>p</i>
Age	1	16	< .001	1	29	< .001
Country of origin	87	1.3	.02	97	2.7	< .001
Gender	1	4.9	.03	1	1.8	.19
Years of education	1	3.6	.06	1	3.4	.07

Table 6. Multiple linear regression analyses of student demographics versus backjumps for certificate earners. We report ANOVA F statistics, degrees of freedom (df), and *p*-values. Age has statistically significant effects in all four courses, country in three courses, and gender in two.

centage breakdown of the four kinds of backjumps. Students most frequently backjump from an assessment to a lecture, which shows that they might be opportunistically looking up specific information needed to answer assessment questions. The second most prevalent kind of backjump is between two lectures, potentially when students are re-watching old lectures for more in-depth learning. Such lecture-to-lecture backjumps occur significantly less frequently (30% of total backjumps averaged across four courses) than assessment-to-lecture backjumps (54% across four courses) (independent t-test, $T_{4,6} = 5.2, p < .01$).

Figure 5 also shows that assessment-to-lecture backjumps are less prevalent in the two programming-based courses (6.00x and CS188.1x). One possible explanation is that programming assignments more often require students to apply concepts to brand-new tasks. This kind of knowledge is difficult to directly look up in earlier lectures, thus making assessment-to-lecture backjumps less helpful than in other courses. Perhaps students jumped more frequently to external Web resources (e.g., StackOverflow or programming tutorial websites) when working on those assignments, but the edX servers cannot log such external resource accesses.

Demographics and backjumps

The prevalence of backjumps also varies by age, country, and gender (see the regression analysis summary in Table 6).

Older certificate-earning students backjump more frequently. Figure 6 visualizes backjumps by age groups, aggregated over all four courses. When we investigated types of backjumps, we found that certificate earners above 40 years of age performed 5% to 12% more lecture-to-lecture backjumps than those under age 20; the differences in proportions are statistically significant for all courses except CS188.1x, using a chi-square test for equality of proportions $\chi^2(1) = 149 - 280, p < .001$. Also, those above 40 years old performed 2% to 11% fewer assessment-to-lecture backjumps than those under 20 years old (statistically significant over all four courses with $\chi^2(1) = 8.8 - 226, p < .003$). One possible interpretation is that older students more frequently review prior




Figure 6. The mean number of backjumps per visited sequence for all certificate-earning students in each age group. Error bars represent the standard error of each age group’s mean. In general, older students backjump more frequently to revisit earlier course materials.

lecture content and do not as frequently work backward using the assessment questions as learning goals.

A student’s country of origin has a significant effect on the proportion of backjumps for three of the courses (see Table 6). To investigate further, we conducted a multiple regression to determine whether a country’s student-teacher ratio partly explains the differences in backjumps. In those same three courses, student-teacher ratio had a negative correlation with backjumps, even controlling for other demographics and grades ($b = -.002$ for 6.00x, $b = -.007$ for PH207x, $b = -.01$ for 3.091x, all $p < .01$).

Figure 7 illustrates this negative correlation for the top 30 most represented countries, which comprise 83% of all certificate earners across the four courses. The top country (U.S.) had 2,809 certificate earners, while the 30th most represented one (Kenya) had 86 certificate earners. This figure shows that a higher student-teacher ratio corresponds to fewer backjumps; the effect is most pronounced for the four countries with the highest ratios. For instance, the country where students backjump most frequently is Greece, with a low student-teacher ratio of 10. Students from Greece backjump, on average, 1.21 times per visited sequence. In comparison, students from Kenya (with the highest student-teacher ratio of 47) have a significantly lower number of backjumps, with .83 times per sequence visit (independent t-test, $T_{212} = 6.0, p < .001$).

However, for older certificate earners, country of origin seems to have no bearing on backjumps. When we conducted the same regression shown in Table 6, except only considering certificate earners over 40 years old, the effects of country were non-significant over all courses ($F_{(28-74)} = 0.9 - 1.3$, depending on the course, all $p > .05$). Older students are more likely self-selected to be independent, self-directed learners, so they might not conform to the general trends of the educational systems in their home countries.

Finally, Table 6 shows a significant effect of gender on backjumps for two of the courses. In both of those courses, men




Figure 7. The mean number of backjumps per visited sequence for all certificate-earning students in the 30 countries with the most students across all four courses. Error bars represent the standard error of each country’s mean. Students from countries with higher student-teacher ratios progress through the course more linearly with fewer backjumps.

performed fewer backjumps than women: In 6.00x, men averaged .97 backjumps per sequence visit, versus 1.13 for women (independent t-test, $T_{6.4} = 858, p < .001$). And in PH207x, men averaged .96 backjumps per sequence visit, versus 1.09 for women ($T_{7.0} = 4235, p < .001$).

Digital textbook usage

The results in the prior sub-section indicate that older students and those from countries with a low student-teacher ratio perform more backjumps. One possible explanation is that this behavior is emblematic of more independent learning and their preferences for non-linear navigation.

This trend is further supported by differences in the number of textbook events between demographic subgroups. Textbooks events are another type of non-linear navigation, since the digital textbook for each course is located outside of the main flow of the course materials. Again, a multiple regression analysis showed that older certificate earners access the textbook more frequently ($b = .00003 - .0001, p < .001$ for all courses). For instance, students over 40 accessed the textbook 27% more frequently than those under 40, normalizing for each student’s total number of events. However, the student-teacher ratio does not have a statistically significant independent effect ($b \sim 0, p > .01$ for all courses).

DISCUSSION AND DESIGN IMPLICATIONS

While we found that most students employ non-linear navigation strategies through MOOCs, we observed differences between demographics, most notably age and country of origin. This section discusses our findings and offers some design recommendations for MOOC platform creators.

General Navigation Strategy

Our findings suggest that most students navigate through the learning content in a non-linear way. Their behavior aligns

with those that researchers have previously called “explorers” [17] in that they do not necessarily follow a given path. It also suggests that although xMOOCs seemingly impose a linear structure [9], students apply characteristics of field-independent learners by defining their own learning paths.

One of the most surprising findings for us was that certificate earners in the four edX courses on average do not visit 22% of course learning sequences. Thus, even though they passed the course, students never even saw a sizable fraction of the course content.

One likely cause for this coverage gap is “open book” assessments, which, once released on the course website, can be viewed along with the rest of the course’s contents. Our analysis of students’ backjumps confirms this assumption: Certificate earners access older content more often after viewing an assessment than after viewing another lecture sequence.

We have two possible explanations for this behavior: First, students might be opportunistically looking up specific information that is needed to answer the assessment questions, which suggests that their motivation is to receive a certificate rather than to advance their knowledge more generally. We saw this assumption further supported by the finding that there are fewer assessment-to-lecture backjumps in the two computer programming courses (6.00x and CS188.1x). Since programming requires applying learned concepts to novel kinds of problems, it is harder for students to find answers to assessment questions by looking at prior lectures. Second, students might be relying on assessments to provide detailed learning goals for the course. If that is the case, then instructors should emphasize learning goals more explicitly throughout the course to help students understand expectations.

These findings also show what distinguishes in-person courses from MOOCs: In conventional classrooms, students at least get exposed to most of the lecture materials. In K12, attendance is usually mandatory, and in universities, it is encouraged. In addition, even when assessments follow an “open book” style, students would likely need to attend lectures before seeing the assessment questions. In contrast, the independence provided by MOOCs means that students can opportunistically work backward from the assessments to the lectures in order to receive a minimum passing grade. The navigation strategies we observed suggest that at least some students try to minimize the effort to earn a certificate.

Design Implications:

One way to spur greater engagement is to replace the current certificates, which only indicate pass/fail, with richer ones indicating grades in combination with additional measures of engagement, such as participation in discussion forums, peer ratings, or time spent solving assessment questions. Such a reward mechanism could motivate students to cover more of the learning content, which, as we found, is positively correlated with grades.

Furthermore, we suspect that certificates from programming courses are more indicative of a student’s skills. In general, assessments requiring an application of previously learned concepts to new problems, such as those necessary in pro-

gramming courses, might therefore be more suitable to assess a student's understanding in MOOCs than those involving more rote forms of learning, such as recall of what was covered in lecture.

Demographic Characteristics in Navigation Strategies

We found significant differences in the navigation strategies of certificate earners depending on their demographics.

Most notably, older certificate earners cover more course materials and repeat more lecture sequences than younger students. This behavior suggests that they follow non-linear, self-defined learning paths, indicative of a field-independent learning style [22]. In comparison to younger students, they also performed more lecture-to-lecture backjumps, and fewer assessment-to-lecture backjumps, indicating that their learning is less driven by the assessment questions. The assumption that older students use more non-linear and self-motivated learning strategies is also supported by the fact that they accessed the course digital textbook – which is optional reading – more frequently than younger students.

Independent of age, we also found that certificate earners from countries with lower student-teacher ratios (e.g., the U.S. and many European countries) cover more content than those who are presumably accustomed to larger classrooms and a more teacher-centered education (e.g., India and Kenya). Students from countries with a high student-teacher ratio not only cover fewer lecture sequences, but they also proceed through the learning materials more linearly and with fewer backjumps than students from low student-teacher ratio countries. This finding confirms prior research showing that students who are used to mostly teacher-centered educational systems predominantly adhere to the learning style of “observers” [17] and field-dependents [5, 12].

We exclude language problems as a possible reason for the negative correlation between student-teacher ratio and coverage, because English is the official classroom language in several of the countries with a higher student-teacher ratio (e.g., in India, Nigeria, and Kenya after a certain age). Also, some of the countries with a lower student-teacher ratio do not have English as an official language, yet their students cover as much learning content as countries such as the U.S., the U.K., or Australia. An alternative explanation is that students in countries with higher student-teacher ratios are somehow more motivated by earning a certificate with minimal effort.

Design Implications:

The different navigation strategies between demographic groups demonstrate that MOOCs are used by students with varying motivations and needs. We believe that MOOC platforms need to be more flexibly engineered to cater to different learning strategies. In particular, younger students and those from countries with higher student-teacher ratios should receive more explicit learning goals to ensure that they know what is expected from them.

For instance, a progress bar highlighting the most important parts of the course could provide more guidance and reduce the independence required to know what needs to be learned. We can also imagine creating motivational mechanisms to

encourage learners to cover more sequences, such as social comparisons (e.g., “other learners usually spent 45 minutes on this sequence”), or counting coverage and other engagement measures toward the final grade.

LIMITATIONS AND FUTURE WORK

A central limitation of this study is that by just analyzing log data, we cannot directly measure students' true motivations, engagement, intent for enrolling in a MOOC, or their knowledge gained after passing the course. Coverage of learning sequences, as well as data on their certificates and grades, therefore served as proxies in our analyses. A more controlled study, perhaps in a blended learning classroom setting, could reveal finer-grained insights in the future.

Also, the population of students from each country who participated in the first edX batch in Fall 2012, and the current generation of MOOCs in general, are probably not representative of the overall population. Specifically, they are likely to be more technology-savvy and English literate, since our corpus consists of math and science courses taught in English. For future studies, it would be desirable to gather more information about students' demographic backgrounds, in combination with knowledge about their intent for enrolling.

Finally, while we were only able to speculate about the reasons for students' navigation strategies, it would be especially interesting to conduct qualitative follow-ups with various demographic groups. This could shed more light on how much students are strategizing or whether a large portion of their behavior is due to a lack of understanding or motivation. We would like to see our results compared to such future studies.

CONCLUSION

The current generation of Massive Open Online Courses such as those provided by edX (so-called xMOOCs) are usually perceived as enforcing a linear, top-down, instructor-provided structure. The findings in this paper showed that despite the linear structure imposed on students – a chronological ordering of weeks and learning sequences – learners predominantly navigate through xMOOCs in a non-linear way. Analyzing the student log data from 140,546 students who participated in four edX MOOCs, we found that, on average, students skip 22% of the learning sequences entirely and perform a high number of backjumps, most often jumping from assessments back to earlier lectures.

While students tend to ignore the linear structure of xMOOCs, younger students and those from countries with higher student-teacher ratios (e.g., India, Kenya, Pakistan) follow the teacher-provided outline more strictly. Independent of demographic background, though, students use assessment questions as an informal guide for what learning content needs to be covered.

Based on these results, we proposed several design ideas that might motivate students to cover more content. Specifically, MOOCs' reliance on binary pass/fail certificates is detrimental to students' motivations for in-depth learning. If certificates listed grades and richer measures of participation and

engagement, we believe they would be more indicative of a student's true knowledge of the course topic.

The results of this paper support the view that MOOCs do not (yet) supersede traditional universities. We especially believe that more work needs to be done to cater to students' different needs and motivations, and to investigate mechanisms that ensure a more in-depth engagement with learning materials.

ACKNOWLEDGMENTS

Thanks to Rob Rubin and Anant Agarwal at edX for enabling this research, Olga Stroilova for helping with data collection, Quanta Computer for funding Philip's postdoc at MIT, and Krzysztof Gajos for his helpful feedback.

REFERENCES

1. MaxMind GeoIP databases and web services. http://www.maxmind.com/en/geolocation_landing.
2. Bates, T. What's Right and What's Wrong About Coursera-style MOOCs?, 2012. Blog entry retrieved from <http://www.tonybates.ca/2012/08/05/whats-right-and-whats-wrong-about-coursera-style-moocs/>.
3. Breslow, L., Pritchard, D. E., DeBoer, J., Stump, G. S., Ho, A. D., and Seaton, D. T. Studying Learning in the Worldwide Classroom: Research into edX's First MOOC. *Research and Practice in Assessment* 8 (2013).
4. Chen, S. Y., and Ford, N. J. Modelling User Navigation Behaviours in a Hypermedia-based Learning System : An Individual Differences Approach. *Knowledge organization* 25, 3 (1998), 67–78.
5. Cheng, Y. W., Sudweeks, F., Cheng, Y. W., and Sudweeks, F. A Longitudinal Study on the Effect of Hypermedia on Learning Dimensions, Culture and Teaching Evaluation. In *Proc. Cultural Attitudes Towards Technology and Communication* (2012), 146–162.
6. Clarà, M., and Barberà, E. Learning Online: Massive Open Online Courses (MOOCs), Connectivism, and Cultural Psychology. *Distance Education* 34, 1 (2013), 129–136.
7. Coetzee, D., Fox, A., Hearst, M. A., and Hartmann, B. Should Your MOOC Forum Use a Reputation System? In *Proc. CSCW'14*, ACM (2014).
8. E. Truluck, Bradley C. Courtenay, J. Learning Style Preferences Among Older Adults. *Educational Gerontology* 25, 3 (1999), 221–236.
9. Ford, N., and Chen, S. Y. Matching/Mismatching Revisited: An Empirical Study of Learning and Teaching Styles. *British Journal of Educational Technology* 32, 1 (2001), 5–22.
10. Grünewald, F., Meinel, C., Totschnig, M., and Willems, C. Designing MOOCs for the Support of Multiple Learning Styles. In *Scaling up Learning for Sustained Impact*. Springer, 2013, 371–382.
11. Hofstede, G. Cultural Differences in Teaching and Learning. *International Journal of Intercultural Relations* 10 (1986), 301–320.
12. Kennedy, P. Learning Cultures and Learning Styles: Myth-understandings about Adult (Hong Kong) Chinese Learners. *International Journal of Lifelong Education* 21, 5 (2002), 430–445.
13. Kizilcec, R. F. Collaborative Learning in Geographically Distributed and In-person Groups. In *AIED 2013 Workshop on Massive Open Online Courses* (2013).
14. Kizilcec, R. F., Piech, C., and Schneider, E. Deconstructing Disengagement: Analyzing Learner Subpopulations in Massive Open Online Courses. In *Proc. Learning Analytics and Knowledge*, ACM (2013), 170–179.
15. Lee, C., Sudweeks, F., and Cheng, Y. The Role of Unit Evaluation, Learning and Culture Dimensions Related to Students Cognitive Style in Hypermedia Learning. In *Proc. Cultural Attitudes Towards Communication and Technology* (2010).
16. Lee, M. W., Y., C. S., Chrysostomou, K., and Liu, X. Mining Students' Behavior in Web-based Learning Programs. *Expert Systems with Applications* 36, 2 (2009), 3459–3464.
17. Liegle, J. O., and Janicki, T. N. The Effect of Learning Styles on the Navigation Needs of Web-based Learners. *Computers in Human Behavior* 22, 5 (2006), 885–898.
18. McLoughlin, C. E. The Pedagogy of Personalised Learning: Exemplars, MOOCs and Related Learning Theories. In *Proc. EdMedia* (2013).
19. Milligan, C., Littlejohn, A., and Margaryan, A. Patterns of Engagement in Connectivist MOOCs. *MERLOT Journal of Online Learning and Teaching* 9, 2 (2013).
20. Naidu, S. Learning About Learning and Teaching Online. *Distance Education* 34, 1 (2013), 1–3.
21. Pask, G. Styles and Strategies of Learning. *British Journal of Educational Psychology* 46, 2 (1976), 128–148.
22. Reed, W., and Oughton, J. Computer Experience and Interval-based Hypermedia Navigation. *Journal of Research on Computing in Education* 30 (1997), 38–52.
23. UNESCO Institute for Statistics. Pupil-teacher ratio, primary. <http://data.worldbank.org/indicator/SE.PRM.ENRL.TC.ZS>.
24. Witkin, H. A., Moore, C. A., and Goodenough, D. R. Field-Dependent and Field-Independent Cognitive Styles and Their Educational Implications. *Review of Educational Research* 47, 1 (1977).