

Different underlying motivations and abilities predict student versus teacher persistence in an online course

Ross M. Higashi¹  · Christian D. Schunn¹ · Jesse B. Flot²

© Association for Educational Communications and Technology 2017

Abstract Free online courses, including Massively Open Online Courses, have great potential to increase the inclusiveness of education, but suffer from very high course dropout rates. A study of 172 K-12 students and 114 K-12 teachers taking the same free, online, summertime programming course finds that student and teacher populations have different underlying motivational models that predict rates of persistence in the course despite having generally similar motivational levels. Student persistence is predicted by prior programming knowledge, intrinsic interest in the subject matter, and mastery approach goals. By contrast, teacher persistence is similarly predicted by intrinsic interest, but then also by self-identity as a programmer, performance approach goals, and negatively by performance avoidance goals. This sub-population discrepancy in predictive factors is novel, and may be reflective of differing environmental conditions or internal mechanisms between students and teachers. Future design of free choice learning environments can take these factors into account to increase rates of user persistence for different target user populations.

Keywords MOOC · Programming · Students · Teachers · Motivation · Persistence

Free online courses, including Massively Open Online Courses (MOOCs), are an increasingly popular form of online instruction offering for large audiences. They typically offer access to live, video-recorded, or online text instruction with no associated fee, and often cover high-demand topics like introductory sciences or computer programming. Since the cost to offer a course to diverse audience scales only with the costs

✉ Ross M. Higashi
rmh57@pitt.edu

¹ Learning Research and Development Center, University of Pittsburgh, 3939 O'Hara Street, Pittsburgh, PA 15260, USA

² Robotics Academy, Carnegie Mellon University, Ten 40th Street, Pittsburgh, PA 15201, USA

of internet bandwidth, institutions offering free online courses very seldom institute a limit on the number of online consumers of the course material. As a result, there are fewer barriers to entry for those who are interested in attending, especially compared to the high costs of official enrollment in a college-level course at the prestigious institutions which often host the courses. This openness has been argued to provide an excellent opportunity for increasing the inclusiveness of education in many content areas (MacLeod et al. 2015).

One often noted phenomenon, however, is the extremely low rate of persistence of users within these courses (Breslow et al. 2013; Gütl et al. 2014). Persistence is influenced by many factors. Among these predictors, MOOC researchers have tended to focus on demographics, courseware interactions, and reported intent for enrollment (Ho et al. 2014; Nesterko et al. 2014a, b; Reich 2014). Individuals' role relationships to the content, e.g. as teachers versus as students, could play a role in determining persistence as well; teaching status has been examined as a predictor of MOOC performance (DeBoer et al. 2013), though not course persistence. In this study, we examine persistence in an online STEM course using two motivational frameworks drawn from the educational research literature—Achievement Goal Theory and Expectancy-Value Theory—which connect motivation and persistence in learning spaces. We further explored an emergent heterogeneity in motivation–persistence relationships between two different kinds of learners who make use of open courses: students (the typically conceptualized learner) and teachers (who use open courses as professional development for traditional instruction).

Background

Achievement goal theory

Achievement Goal Theory (AGT) originated in the 1980s with the primary distinction between learners with mastery goals and performance goals. Mastery goals are those in which the goal is “to develop one’s competence”. Performance goals, on the other hand, revolve around “demonstrating one’s competence by outperforming peers” (Senko et al. 2011, p. 26). This distinction was further refined through empirical work to add a second dimension of approach versus avoidance (Elliot 1999). Approach goals indicate a positive desire for increased competence (Mastery Approach, or MAP) or outperformance of others (Performance Approach, or PAP). Avoidance goals take the form of a double-negation: Performance Avoidance, or PAV, is the desire to avoid the appearance of incompetence. Mastery Avoidance is the theoretical desire to avoid failure-to-learn or loss of previously held skills, although there remains some debate about its proper characterization (Hulleman and Senko 2010).

Overall, learners who report high levels of Mastery Approach goals tend to display many characteristics and behaviors that are beneficial to learning, including increased persistence when facing difficulty (Elliot et al. 1999). Thus, although not previously connected to open course persistence, the connection of achievement goals to learner persistence more generally is well established. Learners who report high levels of Performance Avoidance goals tend toward unproductive behaviors, such as not seeking help when needed; this would presumably lead to lower persistence when faced with a frustrating situation. Empirically, the relationship between Mastery Approach goals and high persistence is well-documented, Performance Approach generally contributes to

persistence (though not always learning outcomes), and the relationship between Performance Avoidance or Mastery Avoidance and persistence is less straightforward (Elliot et al. 1999; Hulleman and Senko 2010).

The latest research on Achievement Goal Theory has attempted to reconcile numerous subtle distinctions. For instance, if a learner wishes to appear competent, there are questions about whether it is the learner's ultimate *reason* behind wanting to appear competent that matters, or whether it is sufficient that the learner wishes to appear competent (Elliot 1999; Senko 2016). There is some evidence to support both perspectives, and the best path forward may include additional distinctions within the theory. However, doing so comes with drawbacks to the theory's parsimony. Such research is in its early stages, and its full theoretical implications remain unexplored (Senko 2016).

As our study's objectives are to investigate the effects of participant motivation on persistence in MOOCs, we employ the slightly older but better understood three-goal model (mastery approach, performance approach, and performance avoidance), acknowledging that it comes with known limitations. In particular, the omission of mastery avoidance from the model mitigates test fatigue and confusion for younger learners (mastery avoidance items typically employ double-negation wording), but discounts adult learners' wishes to avoid losing skills they formerly possessed. Notably, our instrument also measured *appearance* performance goals (e.g. "It is important to me that I look smart compared to other people studying programming") and not *normative* performance goals (e.g. wanting to outperform peers), which are known to have somewhat different predictive effects (Elliot and McGregor 2001; Hulleman and Senko 2010).

Expectancy value theory

A second major theory connecting motivation to actions like persistence is Expectancy Value Theory (EVT), so named for the major constructs it proposes as the key variables in a decision-making process resembling a rational utility calculation (Ajzen and Fishbein 1977; Fishbein and Ajzen 1975; Lovett and Anderson 1996; Wigfield and Eccles 2000). Expectancy is a learner's own expectations for success at a given endeavor. Value is the amount of value the learner would place upon a successful outcome. Cost is the amount of perceived "cost" against which the merits of an action are weighed. If, intuitively speaking, Expectancy times Value exceeds Cost, the learner will continue with a course of action (Wigfield and Cambria 2010), for example, persisting on a task or in a course.

As a construct, Expectancy is generally composed of an individual's beliefs about his or her own level of ability and expectations of success. Value includes several component constructs such as how important the individual feels it is to do well on the task (attainment value); how much the individual enjoys doing the activity (intrinsic value); and how useful the individual thinks the consequences of the action, e.g. skills gained, will be (utility value). Interest in a topic is closely related to intrinsic value because accomplishing something that one is interested in is likely to provide an enjoyable experience (Deci and Ryan 1985; Wigfield and Cambria 2010). Notions of self-identity relate to attainment value by being part of the calculus for deciding how important it is to an individual to do well on a task; and expectancy, by contributing to estimates of personal ability (Eccles 2009). An individual's expectancies and values regarding a particular task (e.g. completing a course activity) are determined by a combination of social influences and feedback from outcomes of previous actions (Wigfield and Cambria 2010). While long acknowledged as present, Cost has only recently become an area of

serious investigation by researchers, and as such, is not included in our measures for this study (Flake et al. 2015).

Predictions about persistence in online courses

MOOC users drop out of online courses for a variety of reasons, including life changes, diminished career relevance, excessive academic load, insufficient academic preparation, lack of social or technical support, and violated expectations about the course format (Gütl et al. 2014). Some of these reasons—particularly academic and social support—overlap substantially with the types of issues that achievement motivation studies have addressed. For instance, persistence in the face of frustrating situations is correlated with learners possessing certain types of achievement goals. Expectancy Value Theory suggests cognitive constructs—expectations of success, practical valuation of the outcome, and others—that weigh against the “cost” of frustration to determine willingness to participate. These motivational factors have been shown to predict persistence in classroom contexts (Ainley et al. 2002; Kuh et al. 2008). However, open courses consist only of volunteers, and thus variation in achievement goals may not be the central predictor for learner persistence here.

In addition, free online course environments do differ somewhat from traditional classrooms in that they are commonly characterized by the presence of external role heterogeneity among learners. Past research has emphasized national, socio-economic, age, and prior knowledge heterogeneity (Breslow et al. 2013; DeBoer et al. 2013). Here we draw attention to another heterogeneity of interest: K-12 students and teachers may both be enrolled in a given course as learners. Indeed, a recent study of MOOC participants (Ho et al. 2015) found that teachers can make up a significant proportion of various MOOCs that were originally intended for traditional students, many actively teaching in the same topic area as the course. In many STEM areas, a significant number of teachers suffer from insufficient background knowledge in the topics they teach (Darling-Hammond et al. 2005). Indeed, teachers and non-teachers perform at similar levels in MOOCs (Deboer et al. 2013). Additionally, MOOC course-taking for professional development reasons is particularly common for individuals under financial or geographic constraints (Dillahunt et al. 2014), so the appearance of teachers as learners in MOOCs is perhaps not so surprising.

Nevertheless, the different external roles of teachers and non-teachers may dictate differences in the two groups' relationship to the course content. Teachers, for instance, might feel compelled to learn the material more thoroughly so as not to appear unknowledgeable in front of their students (i.e., a form of Performance Avoidance orientation). Such a difference might not manifest as a difference in course performance or mean levels of specific motivational constructs, but rather as a differential factor in predicting persistence: if most teachers are subject to this pressure, then teachers with more such motivation may persist at a higher rate; meanwhile, students—not as subject to the same pressure—may show a null or negative relationship between Performance Avoidance motivation and persistence. Thus even under the same theoretical frameworks, it is possible that the linkages between motivational constructs and observed persistence may be substantially different.

Thus, the main research questions to be answered in our study are: (1) whether Achievement Goal Theory and Expectancy Value Theory constructs predict learner persistence in a free online course, and (2) whether the set of predictive constructs is the same or different between K-12 students and teachers enrolled as learners in the same course.

Methods

Course context

We conducted a correlational study of user behavior in a free, online summer course in virtual robotics (i.e., learning to program simulated robots within a virtual 3D environment). The use of virtual robots instead of physical robots reduces the cost, simplifies setup issues, and speeds up the learning process (Liu et al. 2013); these are all factors that are likely to be important in improving diverse access in open online courses. In the US, robotics is a topic that occurs in technology education classes, after school programs and clubs, home school weekend workshops, and summer informal learning programs. The beginner-level programming topics (e.g. movement commands, loops) covered in the course were considered common content that both students and teachers would need to know, so only a single, general-audience course section was made available, open to everyone. Students might take this course if they are interested in programming or in robotics, and teachers might take this course if they were planning on teaching programming or robotics in or after school in the coming years.

The course was comprised of four instructional units targeted toward beginner-level programming concepts (e.g., sequences of commands, loops, if and if-else statements), plus a capstone challenge. See Fig. 1 for an example instructional task. Following a common style of online instruction, typical instructional sequences consisted of a brief “follow-along” video demonstrating the use of a robotics or programming concept, followed by multiple-choice review questions, optional exploration tasks, and after a few such sequences, a programming challenge task based around the new concepts.

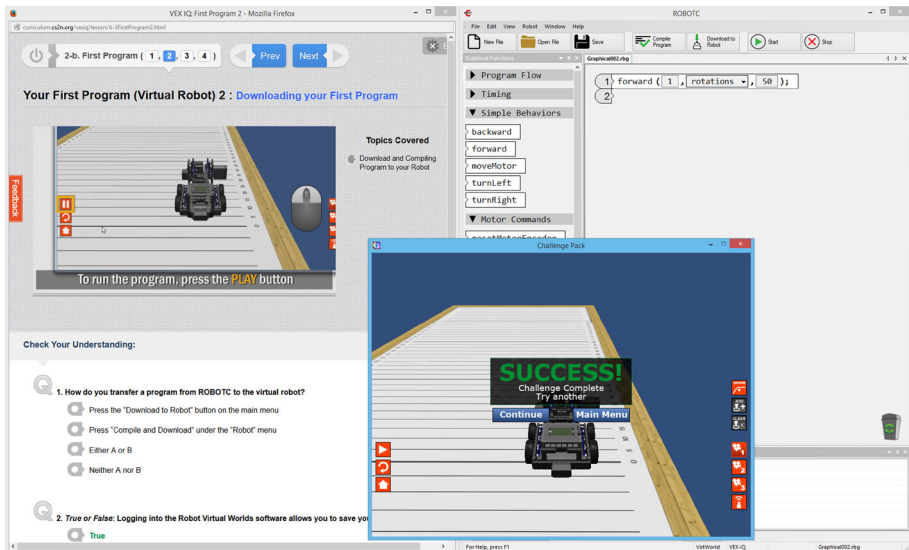


Fig. 1 Screenshot of a typical lesson in the curriculum. Learners watch the video (left, top) and replicate the steps in the programming environment (right). The simulated robot runs the program in a 3D world that mimics a physical setting (center). The lesson page (left, bottom) continues with follow-up questions and additional activities

All instructional materials (instructions for programming challenges, brief open response and multiple choice questions, and interim exams) and research instruments were posted online within a Moodle course. A learner might be expected to complete the full course over 30 h of work. We study their use over a 10-week period during the summer, during which an instructor held live lecture and Q&A sessions at regular intervals. At the end of each instructional unit, learners were required to pass a multiple-choice examination (achieve a score of 70% or higher) on the unit content before being provided with access to the next instructional unit. Users were allowed unlimited attempts without penalty on these examinations. The university-based organization offering the course offered separate “certifications” for students and teachers upon completing the course and achieving a passing score on the final exam.

Participants

Recruitment into the course was done by the organization that created and offered the course. Methods included online advertising through the organization’s websites and e-mail lists, and in-person advertising through conferences and robotics competitions. Participants from a number of different countries signed up for the course, but they were primarily from diverse regions throughout the US; this kind of heterogeneity of learners is common in open online courses. Given our research questions, we selected only K-12 teachers and K-12 students who answered all the pre-survey items (i.e., excluding the small number of adult learners who were not teachers), resulting in a sample of 172 K-12 students and 114 K-12 teachers.

Instruments

In order to gain access to the course materials, each user was required to complete a pre-course activity that served as the key instruments for this research. Participants were not shown the results of the surveys or test, and were allowed to continue on as long as they submitted a response, even if it was left blank. The pre-test contained all of the instruments listed in this section. After each module, participants were asked to also complete the brief attitudinal surveys again.

Demographics

We asked participants to self-identify as one or more of: primary school student, middle school student, high school student, college student, informal educator, formal educator of a non-robotics subject, formal educator of robotics or computer science, homeschool teacher, and/or supervising adult. We later collapsed respondents into simple teacher and student groups, and at that time discarded responses from respondents who checked both teacher and student roles; non-participating adults (e.g. parents who were merely supervising); and college students, who may have fit into either category, but numbered too few to form a third category.

Computer science principles content items

We constructed a set of 14 researcher-designed questions aligned to a programming-focused subset of the recently developed Computer Science Principles (The College Board

2016), the most relevant standards for this topic. To be used sensibly as a pretest of relevant knowledge that could come from many different specific prior programming experiences, these items tested understanding of programming-related skills in non-programming and pseudo-programming contexts (i.e., not knowledge of programming language-specific conventions). Students most commonly began with a high D (just below 70%) and teachers with a high C (just below 80%). Psychometric analyses verified that all items were positively associated with overall performance and no items suffered from ceiling effects.

The example item in Fig. 2 is from a problem sequence involving a bridge that uses an automated traffic signal to control traffic getting on and off a single-lane, single-direction bridge. Vehicle weight is only measurable at the on-ramp, and the off-ramp sensor detects when a vehicle exits, but does not directly measure its weight.

Motivation survey

We asked 21 six-point Likert questions on motivational constructs from both Achievement Goal Theory and Expectancy Value Theory. Items from existing validated surveys (e.g., Elliot and Murayama 2008; Wiebe et al. 2003; Sha et al. 2016) were used directly or adapted slightly to change the topical focus to computer science; the scales that were adapted to computer science were designed to be easily adapted across domains (e.g., “It is important to me that other students think I am good at [insert domain name]”) and have been validated many times across domains and age groups (e.g., Elliot et al. 1999; Elliot and Murayama 2008; Pekrun et al. 2009). Mastery Avoidance items were omitted from our instruments due to concerns about linguistic complexity for younger students (e.g., “One of my goals is to avoid learning less than I could about programming,” could be problematic for elementary school age students), to address concerns about test fatigue, and the weaker prior research connection of this construct to learner persistence.

See Table 1 for the list of constructs and mean (and standard deviations) for each scale, separately for teacher and student subgroups. Teacher and student means were generally similar, except that teachers had slightly higher expectancy of their own success ($M = 5.1$ vs. 4.9), and lower Performance Avoidance motivation ($M = 2.8$ vs. 3.3); see paired t-tests results in Table 1. Standard deviations were similar across students and teachers for each construct, ensuring that differential predictiveness did not result from restricted range problems within one group.

- How would the system know what the “total weight of cars already on the bridge” is?

 - a. Every time a car gets on the bridge, add its weight to the total
 - b. Every time a car gets on the bridge, add its weight to the total, and subtract that amount when it gets off
 - c. Make an estimate by multiplying the last known car’s weight by 100
 - d. It can’t know without an additional weight sensor

Fig. 2 Example test item aligned to CS Principles (4.1.1 Develop an Algorithm)

Table 1 Measures (number of items), sample items, internal reliability, pre-test mean values (with standard deviation), and *t* test significance for groups difference in pre-test means

Measure (# of items)	Sample item	Chronbach's α		Pre-test mean (SD)		Diff?
		Student (n = 172)	Teacher (n = 114)	Student	Teacher	
CS knowledge (14)	See description above	.75	.75	9.2 (3.0)	11.0 (2.6)	p < .01
Programming expectancy (4)	I am sure that I can learn programming	.86	.76	4.9 (1.0)	5.1 (0.7)	p < .05
Intrinsic interest in programming (1)	I love programming	–	–	4.8 (1.2)	4.6 (1.0)	
Identity as programmer (2)	I think like a programmer	.85	.82	3.9 (1.4)	4.0 (1.3)	
Utility value of programming (3)	Programming is a worthwhile and necessary skill	.84	.79	4.9 (1.0)	5.0 (0.9)	
Mastery approach (3)	It is important to me that I thoroughly understand programming	.88	.87	5.1 (1.0)	5.2 (0.8)	
Performance approach (3)	It is important to me that other students think I am good at programming	.88	.82	3.5 (1.4)	3.2 (1.3)	
Performance avoidance (3)	It's important to me that I don't look stupid while learning programming	.84	.89	3.3 (1.4)	2.8 (1.3)	p < .01

All item factor loadings >.3 on the CS Knowledge scale, and >.75 for other scales. A second item for Intrinsic Interest “I would like to learn more about programming” was excluded due to being at-ceiling (M = 5.4 (0.8) for Students and M = 5.4 (0.7) for teachers)

Most scales had acceptable reliabilities as assessed by Chronbach's α levels above .75 (see Table 1). Intrinsic Interest formed a two-item scale with Cronbach's $\alpha = .77$ for students and .68 for teachers. However, the second Intrinsic Interest in Programming item (“I would like to know more about programming”) exhibits signs of a ceiling effect (M = 5.4 on a 6-point scale), so we excluded it from model-building, retaining it only for descriptive purposes. Intrinsic Interest is therefore operationalized by a single-item scale “I love programming” going forward, except where noted.

Analysis

Initial motivation

We began with an analysis of mean motivation levels for students to teachers to examine whether the two differed (and thus have differential predictors of persistence because of greatly differing initial motivational states) and also to generally characterize the motivations of this population.

Persistence

This robotics course, like other free online courses, experienced significant user attrition with each successive unit. Of the $n = 286$ users completing the pre-course activity, only $n = 93$ (32%) reached the end of the second unit, $n = 43$ (15%) reached the end of the third unit, and $n = 10$ (3%) completed the last content unit by the end of the summer.

To have sufficient numbers of completers for statistical analysis, we defined *Persistence* in the course as a binary variable indicating whether a user had completed more than half the available curriculum “units” out of the four units (i.e., 93 of 286), as evidenced by a completed attempt on each unit’s exam. Since a passing score was required on each exam before the next section was made available, users were unable to access later units if they had not successfully completed the prior ones.

We performed separate stepwise logistic regressions to find the factors which achievement goals and which expectancy value factors best explained persistence among students and teachers along with pre-test scores. Factors were sequentially added stepwise if they significantly added predictive power according to the Likelihood Ratio. Exploratory analyses revealed that rate of persistence was generally a linear function of each predictor measure, except Utility, which had a significant quadratic element. However, including a quadratic Utility predictor did not result in a different final best fitting model. Therefore, the presented models involve only linear predictors.

Motivational change

Finally, we examined means and correlations in motivation between the levels reported on the pre-test and those reported at the end of the second unit exam (i.e. matching the cutoff used to operationalize persistence). We refer to values reported on this exam as the “post-test” values relative to our operationalized period for persisting in the course. We examined change across groups between pre and post to assess whether instability in motivation levels accounts for differential predictiveness of the motivational factors across learner groups.

Results and discussion

Who joins an online programming course?

Student learners ($n = 172$) in this course were primarily of middle school (53%) and high school (45%) age. Of the $n = 114$ teachers, 27% were informal educators, 32% were formal educators in non-robotics subjects, 46% formally taught robotics or CS, and 11% were homeschool teachers. Totals sum to more than 100% because educators often teach more than one discipline, or in multiple settings.

As shown in Table 1, users in both the teacher and student groups had relatively high prior intrinsic motivation to participate in the activity. Both teachers and students have a say in the selection of activities in which to participate, especially during the summer. Interest in learning the material, therefore, likely serves as a first-order filter, explaining the overall high value of “I would like to learn more about programming” (the excluded at-ceiling item, with a mean of 5.4 on a 6-point scale). Mean scores for the broader statement “I love programming”, while still in the “Agree” range, were lower (mean of 4.7 on a

6-point scale). The sample group also had high levels of Mastery Approach motivation (mean >5 on a 6-point scale), suggesting that the types of users who are attracted to a MOOC have an existing interest not just in the material, but in mastering it for its own sake. Performance Approach and Performance Avoidance motivation were moderate but more variable among individuals.

Who persists in the online course?

The first stepwise logistic regression analysis of persistence tested the three Achievement Goal Theory factors plus pre-test score (see upper left of Table 2). The procedure resulted in models with different significant predictors for students and teachers. The resulting student model included Mastery Approach motivation and pre-test score. The values shown in the Table are likelihood ratios. For example, a value of 2.17 means a student is 2.17 times as likely to persist for every one-point increase on the 1-to-6 Mastery Approach scale. The effect size for pre-test score is smaller (i.e., it is the equivalent of an 11% increase in persistence for every 10% increase in pre-test score). Other factors did not add significantly above those two included factors.

In contrast to the best student model, the best-fitting teacher model included Performance Approach motivation as a positive factor, and Performance Avoidance motivation as a negative factor (i.e., a completely different set of two predictors). These effects were quite large: every one-point increase in Performance Approach produces an almost 2.5 times increases in persistence, and every *decrease* in Performance Avoidance produced an almost 2 times increase in persistence [note that likelihood ratios below 1 reflect decreasing odds; a likelihood ratio of 0.53 corresponds to a 1.9 times decrease (1/0.53)].

A second set of stepwise logistic regressions with the Expectancy Value factors plus pre-test score once again produced different student and teacher models (see the lower left of Table 2). The resulting student model included Intrinsic Interest and pre-test score. Perhaps because it relied on only a single item, the effect size of the prediction was not

Table 2 Exp(β) (likelihood ratio) for factors in the best and forced logistic regression models of pre-scores predicting learner persistence separately for students and teacher groups

	Exp(β), best model		Exp(β), forced models	
	Students	Teachers	Students	Teachers
Achievement goal scales				
Pre-test score	1.16*			0.99 (ns)
Mastery approach	2.17**			1.89*
Performance approach		2.44**	0.76 (ns)	
Performance avoidance		0.53*	1.24 (ns)	
Expectancy value scales				
Pre-test score	1.17*			0.96 (ns)
Intrinsic interest	1.62**	1.94*	1.71**	2.94**
Identity		1.73*	0.97 (ns)	

Note Pre-Test Score is on a 14-point scale; all others are on a 6-point Likert scale

* $p < .05$, ** $p < .01$, ns $p > .05$ for forced factor

quite as large as the predictiveness of Mastery Approach, but was still an important predictor. The other factors did not add significantly above those included factors. The resulting teacher model included Identity and Intrinsic Interest. The effect sizes of these predictors were large.

To confirm that the student and teacher models were in fact different rather than nearly-as-good for both groups, we ran both groups of logistic regressions again, forcing the opposite group's model on each data set. In the Achievement Goal analysis (see the upper right of Table 2), Mastery Approach motivation was able to explain some of the variance in teacher persistence at a marginal level of significance, but pre-test score was not significant. Neither Performance Approach nor Performance Avoidance were significant for students. In other words, the best teacher model did not fit the students at all, and the best student model only half fit the teachers.

In the Expectancy Value analysis (see the lower right of Table 2), Intrinsic Interest was already present in both models, and remained significant. However, pre-test score was not significant for teachers, and Identity was not significant for students. In other words, the obtained predictive Expectancy Value Theory models for each group were only predictors for that specific group.

Figure 3 presents the separate relationship of each predictor to persistence visually for teachers and students, to examine whether there are similar trends in relationships that are perhaps only statistically significant in one group. For each predictor variable, we divided the variable range, and grouped responses in these bins to compute a mean level of persistence for each bin. The y-axis shows the mean persistence for learners within a given predictor bin, and the x-axis show the included ranges for each bin. The number of bins was adjusted to reflect the distribution of pre-test levels so that no bin had very few responses, but similar results are obtained from using other bin ranges. A best-fitting linear regression line (dotted line) for each group is added to each graph. For only the interest item are the trends in both groups exactly the same (as the logistic regressions also found), and only the Mastery Approach constructs shows the same directional pattern but with different degrees of influence across both groups (as the logistic regressions also found). For the other four motivational constructs, persistence only related to the motivational factor for one group. For example, persistence goes up in students as a function of pre-test score, but the trend is actually slightly negative in teachers.

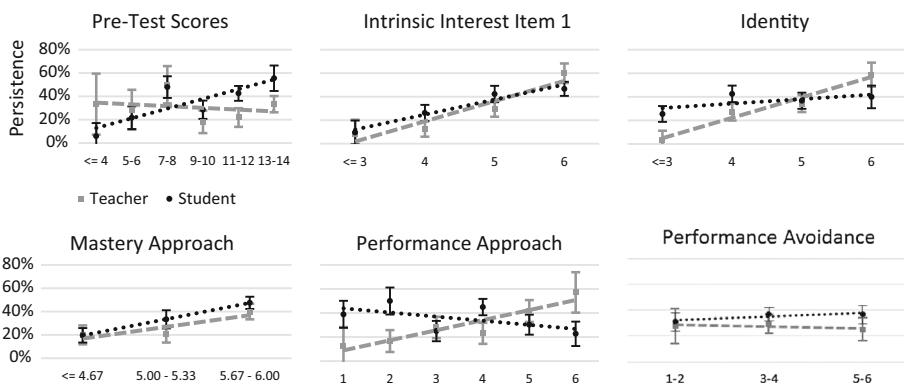


Fig. 3 Separately for teachers (gray squares) and students (black circles), mean persistence (and SE bars) for binned ranges on the six pre-test measures that were significant predictors of persistence for at least one group

Follow-up analyses of motivational change

It is possible learner goals were dynamically changing during the course in uneven ways across groups. For example, perhaps Performance Approach goals was not a significant predictor for students because many students changed the extent which they had these goals during the learning (i.e., the pre-test did not measure stable characteristics). To examine this possibility, we conducted two complementary follow-up analyses on the ($n = 33$) teachers and ($n = 60$) students who had completed the surveys at pre and post: one analysis examined whether there were differential pre-post mean shifts in construct values for teachers versus students, and a second analysis tested for differential relative order stability (or test–retest reliability) within each group.

A two-way mixed ANOVA on teachers and students looked for overall pre-post changes in mean predictor values (see Table 3 for pre and post means). There was no significant shift in most predictors in the model, except for intrinsic interest, which increased for both groups ($p = .055$ for teachers, $p < .001$ for students). However, there were no statistically significant interactions for the constructs included in the best fitting regression models, indicating that teacher and student values shifted in similar ways for all constructs in the model. Expectancy, which was not part of the model, did show a marginally significant ($p = .057$) interaction between time and group, in addition to an overall pre-post shift ($p = .03$).

An examination of the pre-post correlations revealed that most motivational constructs were very stable ($r \approx .7$) (see Table 3), indicating that individuals with low pre-scores on a given measure tended to also have low post-scores, and those with higher pre-scores tended to have higher post-scores. Most importantly, the lowest stabilities were still fairly high (above .4) and predictor stability did not correspond with predictor significance—in particular, lower stability did not prevent predictors from being included in the model: the two measures with the lowest pre-post correlations for teachers—Performance Avoidance motivation ($r = .54$) and Identity as Programmer ($r = .44$)—were both significant in the teacher model. This rules out predictor instability as an alternative explanation for which factors were significant predictors.

Finally, a brief analysis was done to explore the alternate hypothesis that our observed results were the product of advanced students “skipping ahead” and taking the quiz without completing the chapter activities. In other words, did students with high pre-test

Table 3 Test-retest mean scores and Pearson correlation coefficients for motivation construct values before versus after two course modules

	Teachers ($n = 33$)			Students ($n = 60$)		
	Pre-module mean	Post-module mean	Pre-post correlation	Pre-module mean	Post-module mean	Pre-post correlation
Expectancy	5.35	5.36	.74	5.15	5.31	.69
Interest	5.18	5.12	.81	5.20	5.28	.63
MAP	5.37	5.21	.69	5.47	5.42	.75
PAP	3.76	3.82	.69	3.29	3.41	.70
PAV	2.90	3.20	.54	3.30	3.22	.73
Identity	4.73	5.09	.44	4.13	4.73	.74
Utility	5.25	5.23	.69	5.19	5.08	.79

Note All correlation coefficients are significant at an alpha level of $p < .01$

MAP mastery approach, PAP performance approach, PAV performance avoidance

Different underlying motivations and abilities predict...

scores appear to persist but not actually complete all the activities? To test this possibility, we constructed a dummy variable called “skippers” to indicate participants who completed the chapter exam without also completing the final programming challenge in the unit. We then ran a forced logistic regression which found pre-test scores to be a non-significant predictor ($p = .76$) of “skipper” status. When skipping occurs, it is not systematically due to high learner prior knowledge. Thus, pre-test is predicting actual completion of activities (i.e., true persistence).

General discussion

Who persists in an online programming course?

Recall that motivational constructs from both Achievement Goal Theory and Expectancy Value Theory were expected to predict open online course persistence. Here we present a summary of the findings.

Common factors

Intrinsic interest in the subject matter itself (Interest1) predicts both teacher and student motivation to persist in an online course. This is consistent with the EVT view of Intrinsic Interest, where it is a core contributor to Value, and should increase motivation to persist in learning activities. This factor may be especially relevant in the summer learning environment, where a greater variety of competing options like playing outside or going on vacation are viable, compared to the classroom environment, which offers few alternative choices.

Differential factors

There were also different factors, however, predicting student and teacher persistence: (1) Pre-test score predicts persistence for students but not teachers; (2) Performance Approach and Performance Avoidance motivation explain teacher persistence, while student persistence is better modeled using Mastery Approach; and (3) Identity is a significant predictor for persistence in teachers, but not students. Existing motivational theories (and common intuition) would suggest that, since teachers and students were both operating as learners in the course, students and teachers would exhibit the same effects of motivations on behaviors, yet they do not.

That teachers have different needs from students is consistent with what we know about teachers' support needs for classroom practice (Spillane and Zeuli 1999; Mishra and Koehler 2006). But here we saw such differences emerge in a new location: in persistence within the training itself. It is likely that different amounts and combinations of external pressures are present for students and teachers, i.e. the differences in motivational effects are attributable to environmental factors. For example, the demand for teachers to appear knowledgeable about subject material they will later teach might drive them to persist through the learning activities, whereas for students the learning itself may be the primary goal of such an experience (i.e., a difference in the importance of extrinsic versus intrinsic goals).

Alternatively, it could be that greater levels of maturity allow teachers to cope with frustration differently. That is, the differences are potentially due to differences in internal

mechanisms for motivation and the way in which motivation regulates action. Modern theories of self-regulation in learning emphasize the combined importance of motivational and metacognitive factors in guiding learning (Azevedo 2005; Pintrich and DeGroot 1990; Zimmerman and Schunk 2001).

A third possibility, that achievement goals are dynamically changing during the course in ways that are different for students and teachers was ruled out by a follow-up analysis: while there were small overall increases in intrinsic interest for both groups, there were no significant interactions between teacher/student group and time for any of the constructs in the model, and individual pre-post correlations were high.

Finally, there was no evidence to support the alternative explanation that our observed effects of prior knowledge for students were the product of advanced students “skipping ahead” and taking the quiz without completing the chapter activities. While there were some students who skipped ahead, as is often seen in online courses, this behavior was not associated with pre-test knowledge.

Implications for theory

The differentiation in models we observed has implications for theories of motivation in lifelong learning. The finding is not that students and teachers have different *amounts* (i.e., mean values) of motivation, it is that different *types* matter for each group. Current Achievement Goal Theory does not explain how two sub-populations of learners with the same overall levels of Mastery Approach, Performance Approach, and Performance Avoidance motivation can support two different models of persistence. Among students, Mastery Approach goals were a key differentiator between those who persisted in the course, and those who did not. Performance Approach and Avoidance were not predictors for students, even when forced into the model. Teacher persistence, on the other hand, was most strongly predicted by exactly those factors—their levels of Performance Approach and Performance Avoidance motivation were lower on average, but teacher-to-teacher differences more strongly predicted persistence. The effects of Mastery Approach were still present: when forced into the model, Mastery Approach motivation was only a marginally significant predictor.

In the past, the predictiveness of Performance Approach motivation has varied across combinations of learning content and contexts rather than within studies across populations learning the same content. However, in a meta-analysis conducted by Hulleman et al. (2010), students' age (roughly indicated by their grade level in school) was not a moderating factor in the relationship between goal orientations and performance outcomes. That is, the ways in which Mastery Approach, Performance Approach, or Performance Avoidance goal orientations predicted student outcomes were the same even at different grade levels (on average across studies). Of course, it may be that there are qualitative differences once learners reach adulthood that have not been studied sufficiently. The current study suggests there are real differences in the predictive behavior of the different AGT factors in student versus teacher populations. AGT needs to be expanded to explain this type of moderation relationship, whether by external factors (e.g., reasons for task completion) or internal factors (e.g., different regulatory abilities), or a differential stability of achievement goals within student versus teacher groups. Further, it is possible that subgroup effects like the student–teacher differences seen here could explain some divergent findings in past research as well.

Expectancy Value Theory similarly must explain the observed moderation of persistence predictors across groups. At a higher level, EVT predicts that some kind of value is

required for persistence, and both identity and intrinsic interest are kinds of value. But EVT does not currently have a strong accounting for why different Value terms are important predictors between students and teachers: Intrinsic Interest is strong in both groups, but Identity is only a significant predictor for teachers. Perhaps teachers, as older individuals, have a more established sense of self-identity, and use it more effectively to overcome frustration. If this is indeed happening, however, it means that Identity is not functioning as a component of the Value term; it is instead dampening the Cost function.

That pre-test scores *per se* predict student persistence is unexpected. Since participants were not shown their pre-test scores and because expectancy scores were not predictive, it means that the effect is between *actual* prior knowledge and persistence rather than between *perceived* prior knowledge and persistence. Actual knowledge is expected to directly affect learning (Durik et al. 2006) and presumably perceived and actual knowledge are correlated. Indirectly, perceived knowledge can be related to persistence. Because the effect should be indirect (i.e., perceived knowledge affects persistence choices through Performance Approach/Avoidance and Expectancy), actual ability should not normally predict student choices when controlling for achievement goals and expectancy. It could be that pre-test scores are genuinely indicative of a generalizable proficiency in programming-related thinking (i.e. Computational Thinking skills) which enables students to proceed more quickly and with less effort, lowering the Cost function under EVT. Both AGT and EVT might simply interpret such an effect as simply indicating that more proficient students have less frustration to overcome. In either case, the accounting does not explain the fact that this relationship holds for students but not teachers. As before, it suggests the presence of differences in environmental, internal, or goal stability factors.

Implications for practice

Courses like the one we observed are both free of charge and completely open with regard to enrollment, and are thus likely to attract users from more diverse demographic categories than traditional classes. This is favorable from a broader computer science education goal of increasing inclusiveness in the field of Computer Science overall (Smith 2016). However, as logical expectations and our results support, this leads to both a more diverse set of needs, and, if these needs are not met, a greater rate of user attrition. In order to reduce the number of learners who turn away from the course before completing it, appropriate supports must be put in place that recognize and perhaps even take advantage of the non-uniformity of the learner population.

We have now identified a new, important demographic line—students versus teachers—along which learners divide with respect to the ways in which their goal structures affect their likelihood to remain the course. For students, having a greater prior knowledge was associated with persistence. To better support students in these courses, perhaps the pre-test measure could be used to provide additional scaffolds for students with weaker pre-test scores. For example, these students might be directed to additional, easier, challenges. It is interesting, however, that the teachers did not appear to need such additional supports, even when they were at similarly low pre-test levels.

Teachers with a stronger sense of identity to the learning domain persist at a higher rate; while our current findings are only correlational, it now seems reasonable that an intervention designed to increase teachers' sense of programmer identity could improve their rate of persistence (e.g., having teachers write a short essay that asks them to describe how programming fits with their life goals), while doing so with students is less likely to be useful. That some kind of discrepancy in learning occurs between students and teachers is

not new; there are policy and environmental conditions affecting teachers that simply do not apply to students (e.g., credentialing obligations), and it has long been known that there are unique types of knowledge (e.g., Technological Pedagogical Content Knowledge; Mishra and Koehler 2006) which teachers must acquire in order to teach effectively. Professional Development programs that support teachers' mastery of the underlying learning domains have had to take these factors into account in the past; since open courses will include teachers as part of their learner group, they too may need to start making teacher-focused accommodations. For example, since our study indicates that identity as a programmer is positively associated with persistence for teachers (and not harmful to students), an online programming course might deliberately begin calling its participants "programmers" and reinforcing learner self-perceptions of themselves as such. The success of such an intervention could, in turn, provide evidence about whether this relationship is causal in nature.

There are almost certainly more such divisions of learner types within larger MOOCs that could shape the role of motivation in persistence and learning. However, we have also seen that learners who come to the course have one thing in common: an initial intrinsic interest in learning the material. A wisely-designed course might seek to harness this common element, immediately reaching attendees across demographic boundaries. For example, early materials could focus on especially exciting content and allow additional learner choice, two methods that generally improve domain interest and persistence (e.g., Hidi and Renninger 2006; Pelletier et al. 2001). Course designers in open, online learning environments must therefore take the time to understand what kinds of learners are present within their programs, then leverage both universal and group-specific supports to reduce learner attrition in each.

Limitations and future directions

The course we studied bore many resemblances to traditional MOOCs, including: (1) free enrollment and provision of curriculum; (2) opportunities to interact with course staff and other learners; (3) a mix of pre-recorded, live, and interactive materials; and (4) recognition for completion. However, the course size was modest by MOOC standards ($n = 286$), and lacked the resolution of instrumentation (e.g. the "clickstream") that mainstream MOOC platforms provide. Appropriate caution should therefore be taken when generalizing from these results to MOOCs of larger scale. At the same time, there are many smaller scale open courses offered by many different organizations, and the results obtained here will be especially relevant to those kinds of courses.

MOOCs also differ with regard to internal policies around progress and completion. The course we studied used a strict form of progress enforcement that required learners to pass a quiz before proceeding to the next content unit; not all online courses do so. This "gating" mechanism might have been expected to increase attrition overall and strengthen the relationship between persistence behavior and learner knowledge. Yet, in practice, approximately 3% of initial enrollees completed the course, a level comparable to other MOOCs (Gutl et al. 2014). But more importantly, initial learner knowledge predicted persistence for K-12 students but not teachers under these circumstances. If strict progress enforcement created this effect, it would confirm that these two learner groups responded differently to a single designable course feature. If the progress enforcement did not play a role, then it speaks again to the underlying differences between the two groups.

One potential limitation is that our study occurred in a specific, limited timeframe during the summer. Both students and teachers may have different goals during the summer than they do for the rest of the year, as students often engage in more interest-based activities (compared

to normal classrooms), and teachers engage in professional development to prepare for the coming academic year. It is possible that the patterns of motivation we observed are stronger for such elective opportunities than in compulsory attendance environments, or for learning during holidays than for learning that occurs during semesters.

There were also a few theoretical limitations to our results. The omission of a Mastery Avoidance scale from the survey may have omitted an important predictor for adult learners, as Mastery Avoidance—sometimes characterized as the desire to avoid being unknowledgeable or avoid losing existing skills—can be more influential in older adults (Ebner et al. 2006). However, as items seeking to measure Mastery Avoidance are lexically complex due to the multiple negations contained in the construct, the inclusion of such items could simply have introduced a new confound for younger participants and those with lower reading comprehension skills. Additionally, since heterogeneity of motivational predictors was also present among the Expectancy-Value factors, we do not believe the inclusion of Mastery Avoidance would have changed the overall result showing factor \times learner group interactions.

In an analysis of persistence patterns in the HarvardX MOOC, Reich (2014) concludes that persistence rates vary greatly depending upon whether an enrollee intended to complete the course at the outset. While this specific distinction was not present in our data, and we were therefore unable to study or control for it in our design, it is still the case that “intent to complete” (as well as the “reasons for enrolling” reported in Breslow et al. 2013) lack the organized, empirically validated dimensions of construct differentiation upon which our study focused.

As a future research direction, there is a great potential for follow-up investigations into the context surrounding the effects we have observed in each of the two learner groups. Do performance goals matter more to teachers because they are under certain kinds of administrative pressure? This also affords an opportunity to refine survey instruments like our own to account for learner group-specific factors, such as whether teachers are taking the course as a primary form of instruction or as a way to refresh their existing knowledge of programming.

Finally, our study addressed only one type of learner heterogeneity: role differentiation between teachers and students. Many more distinctions could be made in the full online learner population, and we do not endorse teachers versus students as the only important distinction.

Conclusions

Our study examined the ability of motivational constructs from the educational motivation literature to predict user persistence in a MOOC-like online learning environment. We found that such factors were successful in predicting which users would persist, and which would leave, and that different patterns of motivation predicted persistence for different types of users. Persistence among K-12 students was predicted by prior knowledge, mastery approach orientation, and intrinsic interest. Persistence among teachers was predicted by performance approach orientation, performance avoidance orientation, intrinsic interest, and identity. This result underscores the direct relevance of issues of learner diversity to both motivational theory and MOOC design.

Acknowledgements Work on this project was funded by a Grant from the National Science Foundation. We also thank the operators of the online course we studied for ongoing technical support.

Compliance with ethical standards

Conflict of interest Ross Higashi and Jesse Flot were involved in the development of the Robotics Academy curriculum used in the study. Jesse owned stock in Robomatter, Inc. at the time the data was collected, which is the worldwide commercial distributor of curriculum developed at the Robotics Academy. Christian Schunn declares that he has no conflict of interest.

Appendix

Information
[Flag question](#)
[Edit question](#)

Autoautomobile, LLC. is a company developing technologies that will convert normal cars and trucks into self-driving autonomous vehicles.

The questions on this page refer to the development of the Navigation system of the AutoAutomobile technology.

When the user indicates where he or she wants to go, the system conducts a GPS search that finds five possible routes from the vehicle's current location to the requested destination.

Question 4
 Not yet answered
 Marked out of 1.00
[Flag question](#)
[Edit question](#)

Which of the following calculations correctly determines the expected travel time, based on the distance involved and the average speed limit along the route?

Select one:

a. Average speed limit * Total time * Route distance

b. Average speed limit / Route distance

c. Route distance / Average speed limit

d. Route distance * Average speed limit

Question 5
 Not yet answered
 Marked out of 1.00
[Flag question](#)
[Edit question](#)

The system will now know the expected travel time along five different routes. It must then pick the one with the shortest time (smallest value).

Will the following algorithm correctly identify the route with the shortest travel time?

1. Number the five routes 1-5
2. Initially, mark Route #1 the "shortest route" seen so far (in terms of travel time)
3. Then, for each Route from #2 through #5:
 - 3a. Check whether that route is shorter than the existing "shortest route".
 - If it is, mark this as the new "shortest route"
 - 3b. Repeat this process for the next route, until they have all been checked
4. Whichever route is marked the "shortest route" at the end is the actual shortest route.

Select one:

a. Yes, this algorithm is guaranteed to always identify the correct shortest route

b. No, it will always declare Route #1 to be the shortest, even if it is not

c. No, it will always declare the last route to be the shortest, even if it is not

d. The algorithm only works if the routes are already sorted in order of travel time

Question 6
 Not yet answered
 Marked out of 1.00
[Flag question](#)
[Edit question](#)

Which of the following pieces of data could be integrated into the expected-time calculation to help the system calculate a more accurate travel time?

Select one:

a. Amount of traffic along each route

b. Five additional routes to choose from

c. Top speed of the automobile on a test track

d. All of the above could help

[Next](#)

Different underlying motivations and abilities predict...

Information

Flag question

Edit question

The questions on this page refer to the development of the Traffic Safety system of the AutoAutomobile technology.

The autonomous vehicle's collision avoidance subsystem works by limiting the vehicle's maximum speed based on road and traffic conditions. In emergency conditions, it will also take control of the vehicle to avoid a collision.

Question 7

Not yet answered

Marked out of 1.00

Flag question

Edit question

Which of the following expressions could complete the pseudocode line to detect an obstacle too close directly in front of the vehicle?
if (_____), slow down.

Select one:

- a. Distance to nearest object in front < Safety threshold distance
- b. Distance to nearest object in front > Safety threshold distance
- c. Distance to nearest object in front - Safety threshold distance == TRUE
- d. Distance to nearest object in front + Safety threshold distance > 1

Question 8

Not yet answered

Marked out of 1.00

Flag question

Edit question

The following emergency code monitors the vehicle's surroundings, and takes control of the vehicle if a collision is about to occur.

Is the algorithm reasonable?

1. If there is a vehicle dangerously close on the left side:
Steer to the right in order to avoid a collision
2. If not, but there is a vehicle dangerously close on the right side:
Steer to the left in order to avoid a collision
3. If not, but there are vehicles too close on both the left and right sides:
Hit the brakes but do not steer to either side

Select one:

- a. Yes, this algorithm is ready to be written in code
- b. No, the steering actions turn the wrong way
- c. No, Step #3 cannot be reached, because Step #1 would trigger, preventing #2 and #3 from running
- d. No, you cannot have "if" and "if not" decisions in the same program

Question 9

Not yet answered

Marked out of 1.00

Flag question

Edit question

The vehicle safety algorithm calculates multiple values for safe speeds, based on different factors. For example, a car driving in the rain, on a curved road, in traffic might calculate the following:

- Based on current weather conditions, the vehicle should travel no faster than 20 m/s.
- Based on the curve of the road, the vehicle should travel no faster than 15 m/s.
- Based on the distance and speed of the car in front, the vehicle should travel no faster than 10 m/s.

To what speed should the safety system limit the vehicle, and why?

Select one:

- a. 20 m/s because it is the highest value
- b. 15 m/s because it is the average value
- c. 15 m/s because staying on the road is the most important thing
- d. 10 m/s because it is the lowest value

Next

Information

[Flag question](#)[Edit question](#)

The questions on this page refer to the deployment and future development of the AutoAutomobile technology.

Question 10

Not yet answered

Marked out of 1.00

[Flag question](#)[Edit question](#)

The makers of the autonomous driving technology want it to be usable on more than one model of vehicle with only minor adjustments. Is that possible?

Select one:

- a. Yes, if it works on one car, it will work on every car
- b. Yes, but only if the program treats important qualities of the vehicle as variable parameters
- c. No, the technology can only work on the same line of vehicles
- d. No, the technology can only work on the single vehicle it is developed on

Question 11

Not yet answered

Marked out of 1.00

[Flag question](#)[Edit question](#)

To produce a commercially viable product, the developers of the autonomous vehicle technology should plan to:

Select one:

- a. Develop the individual parts of the algorithm, test them, and combine them over time
- b. Develop the entire product as a single piece with no sub-components
- c. Write the program and deploy it immediately without wasting time on testing
- d. Take the time to write code that is 100% perfect and therefore not in need of testing or maintenance

Question 12

Not yet answered

Marked out of 1.00

[Flag question](#)[Edit question](#)

Which of the following is an important type of documentation for the product?

Select one:

- a. Comments in the source code files that explain what important sections and lines of code do
- b. High-level design documentation that explains the way different parts of the system interact
- c. Non-technical documentation of product features written so a user of the product can understand how to operate it
- d. All of the above are important

[Next](#)


Different underlying motivations and abilities predict...

Information
 Flag question
 Edit question

A civil engineer is developing a system that uses automated traffic signals to ensure that the total weight of vehicles on a (single-lane, single-direction) bridge never exceeds a certain amount. A weight sensor measures the weight of each vehicle BEFORE it gets onto the bridge, and a second (simpler) sensor detects when the vehicle reaches the other side and gets off the bridge.

Question 13
 Not yet answered
 Marked out of 1.00
 Flag question
 Edit question

Cars on a bridge are shown on the diagram below, along with their weights and the overall weight limit.



The large truck on the left is waiting to get on, and the maximum weight rating of the bridge is 25,000 lbs. When should the light turn green?

- Select one:
- a. After the next car has gotten off
 - b. After the next two cars have gotten off
 - c. After the next three cars have gotten off
 - d. After all four cars have gotten off

Question 14
 Not yet answered
 Marked out of 1.00
 Flag question
 Edit question

The system should give a car the green light to get onto the bridge as soon as which of the following is true?

- Select one:
- a. There is enough space in front of the car that it will not collide with the car in front of it
 - b. The additional weight of the new car does not cause the total weight on the bridge to exceed the safe maximum
 - c. The weight of the new car does not exceed the maximum total weight
 - d. The light has been red for at least 10 seconds

Question 15
 Not yet answered
 Marked out of 1.00
 Flag question
 Edit question

How would the system know what the "total weight of cars already on the bridge" is?

- Select one:
- a. Every time a car gets on a the bridge, add its weight to the total
 - b. Every time a car gets on the bridge, add its weight to the total, and subtract that amount when it gets off
 - c. Make an estimate by multiplying the last known car's weight by 100
 - d. It can't know without an additional weight sensor that measures the weight of all cars on the bridge

Question 16
 Not yet answered
 Marked out of 1.00
 Flag question
 Edit question

The following pseudocode tries to describe how the "total weight of cars" calculation should behave when a car gets on the bridge, but it contains a bug. What is the bug?

When (car gets on the bridge)
 Set *weight_on_bridge* equal to *weight_of_car*

- Select one:
- a. *Weight_on_bridge* is a constant but should be a variable
 - b. *Weight_on_bridge* and *weight_of_car* are different data types and cannot be used together
 - c. *Weight_on_bridge* should be INCREASED by the weight of the new car, not SET EQUAL TO that amount
 - d. This pseudocode describes what the program should do when a car gets off the bridge, not on it

Question 17
 Not yet answered
 Marked out of 1.00
 Flag question
 Edit question

In principle, the algorithm you developed should work for any bridge, unless which of the following is true?

- Select one:
- a. The bridge doesn't have a weight sensor at the beginning
 - b. The bridge doesn't have a sensor at the end to detect when a car gets off
 - c. The bridge has multiple lanes
 - d. Any of the above would require a different algorithm

References

- Ainley, M., Hidi, S., & Berndorff, D. (2002). Interest, learning, and the psychological processes that mediate their relationship. *Journal of Educational Psychology, 94*(3), 545–561.
- Ajzen, I., & Fishbein, M. (1977). Attitude-behavior relations: A theoretical analysis and review of empirical research. *Psychological Bulletin, 84*(5), 888–918.
- Azevedo, R. (2005). Using hypermedia as a metacognitive tool for enhancing student learning? The role of self-regulated learning. *Educational Psychologist, 40*(4), 199–209.
- Breslow, L., Pritchard, D. E., DeBoer, J., Stump, G. S., Ho, A. D., & Seaton, D. T. (2013). Studying learning in the worldwide classroom: Research into edX's first MOOC. *Research & Practice in Assessment, 8*(1), 13–25.
- Darling-Hammond, L., Holtzman, D. J., Gatlin, S. J., & Heilig, J. V. (2005). Does teacher preparation matter? Evidence about teacher certification, teach for America, and teacher effectiveness. *Education policy analysis archives, 13*(42), n42.
- DeBoer, J., Stump, G. S., Seaton, D., & Breslow, L. (2013). Diversity in MOOC students' backgrounds and behaviors in relationship to performance in 6.002x. In the *Proceedings of the Sixth Learning International Networks Consortium Conference*.
- Deci, E. L., & Ryan, R. M. (1985). *Intrinsic motivation and self-determination in human behavior*. New York: Springer.
- Dillahunt, T. R., Wang, B. Z., & Teasley, S. (2014). Democratizing higher education: Exploring MOOC use among those who cannot afford a formal education. *The International Review of Research in Open and Distributed Learning, 15*(5).
- Durik, A. M., Vida, M., & Eccles, J. S. (2006). Task values and ability beliefs as predictors of high school literacy choices: A developmental analysis. *Journal of Educational Psychology, 98*(2), 382–393.
- Ebner, N. C., Freund, A. M., & Baltes, P. B. (2006). Developmental changes in personal goal orientation from young to late adulthood: from striving for gains to maintenance and prevention of losses. *Psychology and Aging, 21*(4), 664.
- Eccles, J. (2009). Who am I and what am I going to do with my life? Personal and collective identities as motivators of action. *Educational Psychologist, 44*(2), 78–89.
- Elliot, A. J. (1999). Approach and avoidance motivation and achievement goals. *Educational psychologist, 34*(3), 169–189.
- Elliot, A. J., & McGregor, H. A. (2001). A 2 × 2 achievement goal framework. *Journal of Personality and Social Psychology, 80*(3), 501.
- Elliot, A. J., McGregor, H. A., & Gable, S. (1999). Achievement goals, study strategies, and exam performance: A mediational analysis. *Journal of Educational Psychology, 91*(3), 549.
- Elliot, A. J., & Murayama, K. (2008). On the measurement of achievement goals: Critique, illustration, and application. *Journal of Educational Psychology, 100*(3), 613–628.
- Fishbein, M., & Ajzen, I. (1975). *Belief, attitude, intention and behavior: An introduction to theory and research*. Reading, MA: Addison-Wesley.
- Flake, J. K., Barron, K. E., Hulleman, C., McCoach, B. D., & Welsh, M. E. (2015). Measuring cost: The forgotten component of expectancy-value theory. *Contemporary Educational Psychology, 41*, 232–244.
- Gütl, C., Rizzardini, R. H., Chang, V., & Morales, M. (2014). Attrition in MOOC: Lessons learned from drop-out students. In *Learning Technology for Education in Cloud. MOOC and Big Data* (pp. 37–48). New York: Springer.
- Hidi, S., & Renninger, K. A. (2006). The four-phase model of interest development. *Educational Psychologist, 41*(2), 111–127.
- Ho, A. D., Chuang, I., Reich, J., Coleman, C., Whitehill, J., Northcutt, C., Williams, J. J., Hansen, J., Lopez, G., & Petersen, R. (2015). HarvardX and MITx: Two years of open online courses (HarvardX Working Paper No. 10). doi:10.2139/ssrn.2586847.
- Ho, A. D., Reich, J., Nesterko, S., Seaton, D. T., Mullaney, T., Waldo, J., & Chuang, I. (2014). HarvardX and MITx: The first year of open online courses (HarvardX and MITx Working Paper No. 1).
- Hulleman, C. S., Schrage, S. M., Bodmann, S. M., & Harackiewicz, J. M. (2010). A meta-analytic review of achievement goal measures: Different labels for the same constructs or different constructs with similar labels? *Psychological Bulletin, 136*(3), 422.
- Hulleman, C. S., & Senko, C. (2010). Up around the bend: Forecasts for achievement goal theory and research in 2020. In *The decade ahead: Theoretical perspectives on motivation and achievement* (Vol. 16, pp. 71–104). Emerald Group Publishing Limited.
- Kuh, G. D., Cruce, T. M., Shoup, R., Kinzie, J., & Gonyea, R. M. (2008). Unmasking the effects of student engagement on first-year college grades and persistence. *The Journal of Higher Education, 79*(5), 540–563.
- Liu, A. S., Schunn, C. D., Flot, J., & Shoop, R. (2013). The role of physicality in rich programming environments. *Computer Science Education, 23*(4), 315–331.

- Lovett, M. C., & Anderson, J. R. (1996). History of success and current context in problem solving: Combined influences on operator selection. *Cognitive Psychology*, *31*(2), 168–217.
- Macleod, H., Haywood, J., Woodgate, A., & Alkhatnai, M. (2015). Emerging patterns in MOOCs: Learners, course designs and directions. *TechTrends*, *59*(1), 56–63.
- Mishra, P., & Koehler, M. (2006). Technological pedagogical content knowledge: A framework for teacher knowledge. *The Teachers College Record*, *108*(6), 1017–1054.
- Nesterko, S. O., Seaton, D. T., Kashin, K., Han, Q., Reich, J., Waldo, J., Chuang I., & Ho, A. D. (2014). Education Levels Composition (HarvardX Insights).
- Nesterko, S. O., Seaton, D. T., Kashin, K., Han, Q., Reich, J., Waldo, J., Chuang I., & Ho, A. D. (2014). World Map of Education Composition (HarvardX Insights).
- Pekrun, R., Elliot, A. J., & Maier, M. A. (2009). Achievement goals and achievement emotions: Testing a model of their joint relations with academic performance. *Journal of Educational Psychology*, *101*(1), 115–135.
- Pelletier, L. G., Fortier, M. S., Vallerand, R. J., & Briere, N. M. (2001). Associations among perceived autonomy support, forms of self-regulation, and persistence: A prospective study. *Motivation and Emotion*, *25*(4), 279–306.
- Pintrich, P. R., & De Groot, E. V. (1990). Motivational and self-regulated learning components of classroom academic performance. *Journal of Educational Psychology*, *82*(1), 33–40.
- Reich, J. (2014). MOOC completion and retention in the context of student intent. EDUCAUSE Review Online.
- Senko, C. (2016). Achievement goal theory. *Handbook of Motivation at School*, 75.
- Senko, C., Hulleman, C. S., & Harackiewicz, J. M. (2011). Achievement goal theory at the crossroads: Old controversies, current challenges, and new directions. *Educational Psychologist*, *46*(1), 26–47.
- Sha, L., Schunn, C. D., Bathgate, M., & Ben-Eliyahu, A. (2016). Families support their children's success in science learning by influencing interest and self-efficacy. *Journal of Research in Science Teaching*, *53*(3), 450–472.
- Smith, M. (2016). *Computer science for all* | whitehouse.gov. Retrieved from <https://obamawhitehouse.archives.gov/blog/2016/01/30/computer-science-all>.
- Spillane, J. P., & Zeuli, J. S. (1999). Reform and teaching: Exploring patterns of practice in the context of national and state mathematics reforms. *Educational Evaluation and Policy Analysis*, *21*(1), 1–27.
- The College Board. (2016). *AP computer science principles*. New York, NY. Retrieved from CS Principles website: <https://secure-media.collegeboard.org/digitalServices/pdf/ap/ap-computer-science-principles-course-and-examdescription.pdf>
- Wiebe, E., Williams, L., Yang, K., & Miller, C. (2003). Computer science attitude survey. *Computer Science*, *14*(25), 0–86.
- Wigfield, A., & Cambria, J. (2010). Expectancy-value theory: Retrospective and prospective. *Advances in Motivation and Achievement*, *16*, 35–70.
- Wigfield, A., & Eccles, J. S. (2000). Expectancy–value theory of achievement motivation. *Contemporary Educational Psychology*, *25*(1), 68–81.
- Zimmerman, B. J., & Schunk, D. H. (Eds.). (2001). *Self-regulated learning and academic achievement: Theoretical perspectives*. New York: Routledge.

Ross Higashi is a Graduate Student Researcher in the Learning Sciences and Policy program at the Learning Research and Development Center of the University of Pittsburgh. He studies digital badges, robotics-based STEM education, and game-based online learning. He previously designed robotics curricula for the Carnegie Mellon Robotics Academy, where he is now retained as a part-time consultant.

Christian Schunn is a Senior Scientist at the Learning Research and Development Center and a Professor of Psychology, Learning Sciences and Policy, and Intelligent Systems at the University of Pittsburgh. He directs a number of research and design projects in science, mathematics, and engineering education. This work includes studying expert engineering and science teams, building innovative technology-supported STEM curricula, and studying factors that influence student and teacher learning and engagement.

Jesse Flot is a software engineer at the Carnegie Mellon Robotics Academy. As the product manager for Robot Virtual Worlds, Jesse leads the planning for simulation and gaming environments designed to teach programming, math, and other STEM related concepts. He is also the lead writer for several Robotics Academy curricula designed to teach robot programming and big ideas in Computer Science.