Introduction to Online Learning at Scale

Online learning at scale enables new learning opportunities for vast numbers of students around the world. This offers the human-computer interaction community at least two significant directions of research: we can invent their experience, and we can study how those students learn. This special issue includes five articles that report on these directions and more.

New interests, jobs, challenges, and hobbies appear throughout our lives. How might we expand lifelong learning opportunities to match them? Online learning at scale promises to provide always-on education that is available lifelong and around the globe. Research labs and classrooms have been creating multimedia learning experiences for decades. Interest has amplified dramatically in recent years, as producing and consuming video and multimedia has become increasingly cheap, fast, and pervasive. Online course providers leverage this for massive scale and massive news coverage; the New York *Times* declared 2012 the year of the MOOC. What has happened since then?

This special issue explores the research frontier of online education at scale. In particular, this special issue's focus on massive scale explores how it transforms, challenges, and enables online learning. Although many building blocks of online education are decades (or, in some cases, millennia) old, the massive scale of participation make for a new and different setting. One major impact of the scale and online setting is that it makes experimentation practical. Consequently, online learning is a great petri dish for learning science because for the first time, we can practically manipulate learning approaches and measure learning outcomes. We get to simultaneously change the world through education and use the platform as a test bed for scientific explorations of individual and social learning.

With scale, infrequent responses become numerous. If 2% of students make a particular error, that's just 2 students in a class of 100—a rate at which teachers probably won't notice. But that same error rate is 200 students in a class of 10,000. At that scale, we can start to identify why those 200 students misunderstood something and how to avoid that in the first place. Tiny signals suddenly become very visible.

This theme of scale will resonate for many readers in the human-computer interaction (HCI) community, as scale has been similarly disruptive and enabling for many areas of HCI. It is our hope that these readers of this special issue will find insights into how to leverage scale and also a domain for applying the insights from their own work.

The articles in this issue create "field notes" from the future to guide innovation in online learning. The articles also analyze the current state of affairs, giving us a sense of where online learning stands at this time. In addition to the individual value of the articles, it is our hope that readers will benefit in two ways from reading them collectively. First, these articles sketch the breadth of the field. For those interested in getting involved or benefiting from this research, the special issue gathers the different types of work under way. Second, the literature relevant to online learning spans several fields. The articles in this issue do an admirable job of synthesizing this disparate prior work and operationalizing fundamental insights for use in online education.

As the social psychologist Kurt Lewin wrote, "If you want to understand something, try and change it." For those of us interested in the cognitive and social processes of learning, online education provides an exciting and powerful petri dish to try and create

DOI: http://dx.doi.org/10.1145/2737794

change for the better. Many fundamental issues in HCI are, at their core, questions about learning. And many key opportunities for improving learning are, at their core, sites for improving HCI.

One key riddle of online education is how to take *inspiration* from the best models we have from face-to-face education—such as direct tutoring (Bloom), mastery learning, and studio critique—without being *shackled* to them. The most exciting pedagogical approaches for online learning at scale go *beyond* being there (Hollan). In particular, many successful approaches treat scale as a *resource*, not just a challenge. For example, perhaps the only thing that automatically scales with the number of students in a class is the number of students in a class. This tautological insight explains part of the power and appeal of peer learning. Furthermore, global classrooms naturally exhibit a cultural and intellectual diversity that can be harnessed as a learning resource.

An important motivation for this special issue is that currently, a lot of exciting work is happening in disparate venues, yet this work is being driven by synergistic forces. For example, one recent sea change is the breadth of availability of networking and computation that makes video-on-demand unremarkable and inexpensive. This opens the door to global, large-scale learning communities. What do we want the one-world classroom to be? Some online education models offer dramatic cost reductions relative to on-land models. Other uses of online education focus more on increased quality than decreased cost.

One grand challenge for online education is to create a personalized, motivating, mastery-learning experience for every student. Students vary in their goals and background. Traditional classrooms require all students to move at the same pace, even though that's too quick for some, too slow for others, and bundled in ways that may not meet learner goals and interests. Several of the articles in this issue make important contributions toward this goal of "mass customization" for learners. If successful, this will have important equity benefits, as students with different backgrounds and motivations can find an entree to learn more. As you read these articles, give some thought to how the advances described in this special issue might change (or not) the educational architectures that we experience today, both "full-time" (like universities) and just-in-time.

For all of this excitement and opportunity, we don't mean to imply that research advances are inevitable or easy. Education research, even with the benefits of scale, remains notoriously difficult. Partially for this reason, early characterizations of MOOCs focused on descriptive statistics and finding correlations in existing datasets. The challenge for the next wave of research is to move to experimentation and measure learning, and the articles in this issue rise to the challenge!

Online education research covers a vast terrain, and the five articles in this issue represent excellent instances of different parts of that space. This sample will give you a flavor of different parts of the field, and some "hooks" such that you can use these articles as guides to help you dive deeper into areas that interest you.

In "Motivation as a Lens for Understanding Online Learners: Introducing and Applying the Online Learner Enrollment Intentions (OLEI) Scale," Kizilcec et al. make a compelling point that we might not really understand the motivations of students who take online classes. They introduce the OLEI survey, which was derived from their analysis of many open-ended responses about why the learners are there. Testing this survey in 14 MOOCs, they found the assumption that learners are primarily motivated to finish lectures and assessments to earn a certificate of completion is unfounded. Instead, only about half the learners in the sample ($N = 71{,}475$) reported enrolling with this as their goal. This article sheds light on the motivations of the students, but one important characteristic of MOOCs is that there is nearly always a huge variety of students in the class. They vary by motivation, but also by background, experience, country of origin, and class-relevant skills.

The article by Glassman et al., "OverCode: Visualizing Variation in Student Solutions to Programming Problems at Scale," gives us insight into the range of solutions that can be generated by students in a MOOC. To give decent feedback to this wide variety of students, we need to understand the variation among these solutions. That variation can be mined as a source of pedagogically valuable examples and used to refine the course autograder by exposing corner cases. OverCode clusters similar solutions together, allowing instructors to filter and cluster solutions based on different criteria and thus gaining insight into what's working and what's not.

Students vary by motivation, participation, background, and expertise. In sophisticated domains, this puts a premium on being able to create automatic feedback systems to help guide students in their studies.

D'Antonio et al. answer an important question in their article, "How Can Automatic Feedback Help Students Construct Automata?" As large online teaching systems tackle increasingly complex, technical domains, it becomes critical to give students access to an online laboratory to explore how things work. The authors report on their experiences teaching computer science students to build finite-state automata. As with many domains, it's now possible to run simulations of the teaching domain and actually test out crazy ideas and play around with the things in the domain. Building such a supportive online environment requires many widgets with domain smarts and intelligence. This is probably a central theme of many learning environments that will be constructed over the next several years. In this article, we see how the leading edge of building smart online domain simulation environments actually operates.

In "Large-Scale Educational Campaigns," Liu et al. point out that it often takes an educational "campaign" to market the content of the class, simultaneously generating awareness of the course while getting students motivated to participate. As is clear from other articles here, the larger the class, the greater the social learning effects that can take place. Whereas universities and schools are essentially their own market, the new large-scale classes are qualitatively different and have yet to have their own "natural" marketplace. This article analyzes three campaigns, showing what works well (and less well) in three very different markets. Their work also highlights the effectiveness (or not) of letting students "test out" of parts of the course to fast-forward themselves through the material and offers other kinds of incentives as well.

Konstan et al. show us the pragmatics of delivering their hybrid MOOC in "Teaching Recommender Systems at Large Scale: Evaluation and Lessons Learned from a Hybrid MOOC." The goal of their MOOC is to teach engineering students how to understand and build recommender systems, such as those that power commercial book and music recommendation systems. This article is essentially their field report and their reflection on teaching the course both as a MOOC and as a live, in-person class. They report on the *actual* behavior of students, and they contrast the live class with the online class, offering insights into what matters (and doesn't matter) in the students' experience. As these kinds of field experiences accumulate, we can start to compare the design and implementation of large online courses across different fields, platforms, and domains. This work is an important first step in capturing the effect of the entire class—preclass, midclass, and postclass.

So enjoy these articles, as it is our hope that they will inspire you to pursue your own research in online learning at scale.

Benjamin B. Bederson
Daniel M. Russell
Scott Klemmer

Guest Editors