Students, Systems, and Interactions: Synthesizing the First Four Years of Learning@Scale and Charting the Future

Sean Kross
UC San Diego
La Jolla, CA, USA
seankross@ucsd.edu

Philip J. Guo
UC San Diego
La Jolla, CA, USA
pg@ucsd.edu

ABSTRACT
We survey all four years of papers published so far at the Learning at Scale conference in order to reflect on the major research areas that have been investigated and to chart possible directions for future study. We classified all 69 full papers so far into three categories: Systems for Learning at Scale, Interactions with Sociotechnical Systems, and Understanding Online Students. Systems papers presented technologies that varied by how much they amplify human effort (e.g., one-to-one, one-to-many, many-to-many). Interaction papers studied both individual and group interactions with learning technologies. Finally, student-centric study papers focused on modeling knowledge and on promoting global access and equity. We conclude by charting future research directions related to topics such as going beyond the MOOC hype cycle, axes of scale for systems, more immersive course experiences, learning on mobile devices, diversity in student personas, students as co-creators, and fostering better social connections amongst students.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords
context/synthesis paper; survey paper; meta-analysis

INTRODUCTION
Although learning technologies have been around since the dawn of the computer age in the 1960s with early systems like PLATO [87], over the past decade the massive increases in computing power and worldwide internet access have enabled these technologies to scale to reach students in almost every country. In response to recent trends, the Learning at Scale conference formed in 2014 as an interdisciplinary venue where researchers from fields such as learning science, computer science, social sciences, and design come together to study and address the sociotechnical challenges of scaling learning.

L@S 2018, June 26–28, 2018, London, United Kingdom

ISBN 978-1-4503-5886-6/18/06...$15.00

DOI: https://doi.org/10.1145/3231644.3231662

Now that this conference has completed four years (it has metaphorically "graduated from college") and published 69 full papers, it is an appropriate time to take a step back to reflect on where we have been as a community and where we might go in the coming years. In this context/synthesis paper, we survey all four years of Learning at Scale papers to reflect on the major research areas that have been studied so far and use those insights to chart possible directions for future study.

Why reflect on only the papers at Learning at Scale instead of expanding our survey to include other related publications? Because this new community was formed in recent years with the explicit goal of addressing scale and bringing together disparate fields instead of having legacy roots entrenched in any one field. Since no single field is fully equipped to address the challenges of designing the future of technology-enabled education for a worldwide audience, we believe that the Learning at Scale publication record represents one of the most representative overviews of the state-of-the-art in this space.

The contributions of this paper are:

- A taxonomy of all 69 full papers from all Learning at Scale meetings so far (2014–2017), which reveals the major clusters of research activity at this conference.
- A summary of representative papers from this taxonomy, which introduces readers to the flavor of research done here.
- Design ideas for future research directions inspired both by these existing papers and by evolving technological trends.

TAXONOMY OVERVIEW
Table 1 shows how we classified all 69 full papers from all four years of Learning at Scale so far (2014–2017). Our approach was to work bottom up in an inductive manner by first reading all of the papers and trying to characterize their primary research contributions. At the highest level, we noticed that nearly all papers contributed either software systems or empirical studies. Out of the study papers, the next most salient distinction was whether researchers were studying how students interacted with existing software systems (e.g., how they watch MOOC videos or use discussion forum software) or studying the properties of online students themselves (e.g., their levels of self-efficacy or knowledge). Thus, we created three top-level categories: SYSTEMS FOR LEARNING AT SCALE, INTERACTIONS WITH SOCIOTECHNICAL SYSTEMS, and UNDERSTANDING ONLINE STUDENTS. We then partitioned each category into a second level of sub-categories.

SYSTEMS FOR LEARNING AT SCALE	
One-to-One Systems	Bayesian Ordinal Peer Grading [67] BayesRank: A Bayesian Approach to Ranked Peer Grading [82] Graders as Meta-Reviewers: Simultaneously Scaling and Improving Expert Evaluation for Large Online Classrooms [35] Improving the Peer Assessment Experience on MOOC Platforms [76] Peer Grading in a Course on Algorithms and Data Structures: Machine Learning Algorithms do not Improve over Simple Baselines [71] PeerStudio: Rapid Peer Feedback Emphasizes Revision and Improves Performance [46] Scaling Expert Feedback: Two Case Studies [34] Scaling Short-answer Grading by Combining Peer Assessment with Algorithmic Scoring [47] Self-evaluation in Advanced Power Searching and Mapping with Google MOOCs [85]
One-to-Many Systems	Divide and Correct: Using Clusters to Grade Short Answers at Scale [4] Gradescope: A Fast, Flexible, and Fair System for Scalable Assessment of Handwritten Work [72] Teaching Students to Recognize and Implement Good Coding Style [83] Writing Reusable Code Feedback at Scale with Mixed-Initiative Program Synthesis [31]
Many-to-Many and Automated Systems	An Automated Grading/Feedback System for 3-View Engineering Drawings using RANSAC [48] An Exploration of Automated Grading of Complex Assignments [25] Autonomously Generating Hints by Inferring Problem Solving Policies [65] AXIS: Generating Explanations at Scale with Learnersourcing and Machine Learning [86] Enabling Real-Time Adaptivity in MOOCs with a Personalized Next-Step Recommendation Framework [64] Fuzz Testing Projects in Massive Courses [73] Hint Systems May Negatively Impact Performance in Educational Games [59] How Mastery Learning Works at Scale [69] Learning is Not a Spectator Sport: Doing is Better than Watching for Learning from a MOOC [39] Mathematical Language Processing: Automatic Grading and Feedback for Open Response Mathematical Questions [50] moocRP: An Open-source Analytics Platform [63] Problems Before Solutions: Automated Problem Clarification at Scale [1]
INTERACTIONS WITH SOCIOTECHNICAL SYSTEMS	
Individual Interactions with Learning Technologies	Demographic Differences in How Students Navigate Through MOOCs [29] Detecting Diligence with Online Behaviors on Intelligent Tutoring Systems [14] Do Performance Trends Suggest Wide-spread Collaborative Cheating on Asynchronous Exams? [8] Effects of In-Video Quizzes on MOOC Lecture Viewing [42] Explaining Student Behavior at Scale: The Influence of Video Complexity on Student Dwelling Time [80] How Video Production Affects Student Engagement: An Empirical Study of MOOC Videos [28] Student Skill and Goal Achievement in the Mapping with Google MOOC [84] Superposter behavior in MOOC forums [32] Understanding In-Video Dropouts and Interaction Peaks in Online Lecture Videos [36] Using Multiple Accounts for Harvesting Solutions in MOOCs [70]
Group Interactions within Online Communities	$1 Conversational Turn Detector: Measuring How Video Conversations Affect Student Learning in Online Classes [74] Addressing Common Analytic Challenges to Randomized Experiments in MOOCs: Attrition and Zero-Inflation [49] Alumni & Tenured Participants in MOOCs: Analysis of Two Years of MOOC Discussion Channel Activity [56] A Playful Game Changer: Fostering Student Retention in Online Education with Social Gamification [43] Blended Learning in Indian Colleges with Massively Empowered Classroom [13] Chatrooms in MOOCs: All Talk and No Action [10] Do Professors Matter? Using an A/B Test to Evaluate the Impact of Instructor Involvement on MOOC Student Outcomes [78] Monitoring MOOCs: Which Information Sources Do Instructors Value? [77] Online Urbanism: Interest-based Subcultures as Drivers of Informal Learning in an Online Community [26] Staggered Versus All-At-Once Content Release in Massive Open Online Courses: Evaluating a Natural Experiment [55] Structure and Messaging Techniques for Online Peer Learning Systems that Increase Stickiness [41] Teaching Recommender Systems at Large Scale: Evaluation and Lessons Learned from a Hybrid MOOC [40] The Role of Social Media in MOOCs: How to Use Social Media to Enhance Student Retention [89]
UNDERSTANDING ONLINE STUDENTS	
Modeling Student Knowledge	A Data-Driven Approach for Inferring Student Proficiency from Game Activity Logs [22] A Visual Approach towards Knowledge Engineering and Understanding How Students Learn in Complex Environments [23] Brain Points: A Deeper Look at a Growth Mindset Incentive Structure for an Educational Game [60] Effective Sampling for Large-Scale Automated Writing Evaluation Systems [18] Epistemic Cognition: A Promising and Necessary Construct for Enriching Large-scale Online Learning Analysis [33] Exploring the Effect of Confusion in Discussion Forums of Massive Open Online Courses [88] Improving Student Modeling Through Partial Credit and Problem Difficulty [61] Learning Transfer: does it take place in MOOCs? [9] Probabilistic Use Cases: Discovering Behavioral Patterns for Predicting Certification [11] Robust Evaluation Matrix: Towards a More Principled Offline Exploration of Instructional Policies [17] The Prediction of Student First Response Using Prerequisite Skills [2] Towards Detecting Wheel-Spinning: Future Failure in Mastery Learning [27] Uncovering Trajectories of Informal Learning in Large Online Communities Of Creators [88]
Promoting Global Access and Equity	Attrition and Achievement Gaps in Online Learning [38] Correlating Skill and Improvement in 2 MOOCs with a Student's Time on Tasks [7] Learning about Learning at Scale: Methodological Challenges and Recommendations [81] Learning to Code in Localized Programming Languages [15] Mobile Devices for Early Literacy Intervention and Research with Global Reach [3] Preventing Keystroke Based Identification in Open Data Sets [52] The Civic Mission of MOOCs: Measuring Engagement across Political Differences in Forums [68] Towards Equal Opportunities in MOOCs: Affirmation Reduces Gender & Social-class Achievement Gaps in China [37]

Table 1. Our taxonomy of all 69 full papers from Learning at Scale 2014, 2015, 2016, and 2017. (Paper order within each category is not significant.)

Note that no single taxonomy can fully capture the nuances of all papers, and we found that some papers had multiple types of contributions; in those cases, we put each under what we deemed as the category that best fit its primary contribution. A few papers made methods or theory contributions, so we put them into the categories that most closely fit their subjects.

Limitations: This taxonomy exclusively includes publications from Learning at Scale, which represents only a part of the technology-enabled learning literature. A more complete analysis might include conferences like Educational Data Mining or Learning Analytics and Knowledge. We believe that the methodology for creating this taxonomy is relatively straightforward, however we recognize that there are more systematic methods that have been used in adjacent fields [57], including the Delphi method for group judgements [54].

In the following sections, we survey a representative sample of papers from each category and sub-category in our taxonomy. Due to space constraints, we cannot summarize all 69 papers, so we picked a subset that embodies the themes of each category. The remaining papers are all cited in Table 1.

SYSTEMS FOR LEARNING AT SCALE

The first set of papers from Table 1 that we synthesize are those where researchers built interactive systems to support learning at scale. These systems are often deployed in MOOCs or large university courses. Although we categorized them by the degree to which they potentially amplify human effort (e.g., one-to-one, one-to-many, and many-to-many systems), we also observed that feedback, grading, and hints were three other common themes that cut across all sub-categories.

One-to-One Systems

The canonical one-to-one system presented at Learning at Scale facilitates *peer feedback*, where each student gives one-to-one asynchronous written feedback to evaluate the work of their fellow students [76]. Researchers found that this strategy approximates the quality of expert feedback [35] and is also beneficial to the student giving the feedback [82].

PeerStudio efficiently recruits students to quickly provide focused feedback [46]. Kulkarni et al. report that feedback that is not quickly received has the same effect as receiving no feedback at all with respect to the student's grade. Students have several motivations for soliciting feedback, including wanting comments about how a skill is expressed in their work (like their grammar), comments on a specific part of the assignment, a "sanity check" before submitting their assignment, or just to hear their classmates' opinions. PeerStudio shows how a human-centered system can organize students so that they can be motivated to strengthen each other's work.

Even in cases when peer feedback is not sufficient for grading assignments, it can still be valuable for enhancing expert feedback. Joyner et al. deployed a peer feedback system as part of Georgia Tech's Online Masters of Science in Computer Science program [35]. Expert graders who were able to see peer feedback while grading provided significantly better feedback to students. In a separate study [34] Joyner examined how Coursera, which provides an industry microcredential certificate, trains high-achieving students to provide expert reviews. The feedback of these experts-in-training is itself reviewed, ensuring that the expert feedback that students will eventually receive is consistent. Training experts appears to be an effective strategy for providing fast feedback, with a review occurring every two minutes, and the median time of 92 minutes from submitting an assignment to receiving a review.

Researchers have also highlighted limitations of peer feedback. Kulkarni et al. observed that multiple peer reviewers are prone to give supportive but erroneous feedback to the same submission [47]. Both Waters et al. [82] and Raman et al. [67] present approaches for steering peer feedback in productive directions by asking students to rank a small number of submissions. Both use Bayesian methods to order these rankings. The BayesRank model developed by Waters et al. optimizes which submissions are ranked by which peers, while Raman et al. precisely quantify the uncertainty of the rankings.

Reflection: Systems can empower direct student-to-student interaction, often the most personalized asset in a course.

One-to-Many Systems

The canonical one-to-many system for learning at scale amplifies the efforts of one individual (usually an expert instructor) to reach a large number of students.

Scaling grading and assignment feedback is a significant challenge in MOOCs and large in-person courses. Several systems presented at Learning at Scale work to ensure that the individual expertise and feedback from a single instructor can be applied to many students. Brooks et al. developed a system for grading short answer questions by clustering responses using standard text classification techniques [4]. They found that clustering responses allowed instructors to complete grading more quickly with no reduction in grading accuracy. Additionally more answers received feedback, and instructors were able to spend more time giving more detailed feedback.

Head et al. developed one-to-many systems for programming assignments: FIXPROPAGATOR and MISTAKEBROWSER [31]. Both aim to cluster feedback on code improvements: FIXPROPAGATOR sources these improvements from corrections that instructors make to students' submitted code, while MISTAKEBROWSER sources improvements from students' fixes to their own code. The generated bug fixes improved grading time, while also giving instructors a better sense of common misconceptions among the students in the class. The task of correcting students' errors with a higher level of personalized feedback is then made more manageable via clustering. Using tractable machine learning techniques improved both the quality of feedback that students get and the instructor's understanding of how students are progressing through a course.

Similarly, AutoStyle provides real time feedback about whether students are using good coding style, offering suggestions tailored to the stylistic faux pas that the student is committing [83]. AutoStyle also uses clustering on student code submissions to identify specific interventions that can be made to improve code style. Wiese et al. find that not only does AutoStyle help improve students' coding style, but it also improves students' ability to recognize good coding style when they see it. AutoStyle is in essence acting as an assistant instructor, with the ability to provide feedback that is temporally and spatially much closer to the student's learning environment.

Reflection: Experts can effectively scale their reach by using data to direct their feedback to clusters of students.

Many-to-Many and Automated Systems

Generalizing beyond one-to-many, AXIS [86] embodies a many-to-many system by letting students submit their own explanations of course concepts; over time it learns which explanations were best understood by students, and eventually the quality of the explanations offered rivaled explanations written by experts. As opposed to one-to-many systems, the main insight demonstrated in a system like AXIS is that multiple student inputs can converge to a consensus, which can then be broadcast to subgroups of other students.

More generally, fully automated systems scale even better since they can provide personalized feedback to each student on demand without waiting for human intervention.

Many such systems try to provide feedback in the way a human grader might provide, but doing so in an automated manner. For instance, Kwon and McMains developed RANSAC, a system for providing nuanced feedback for complex technical drawings [48]. It is difficult to provide feedback for technical drawings because rotations, translations, and the positioning of a student's drawing may actually be correct even though it is not identical or congruent to the correct drawing. RANSAC is able to manipulate technical drawings so that errors in a student's drawing are obvious, even to the extent that RANSAC finds errors that are often missed by human graders.

As another representative example, MLP (Mathematical Language Processing) by Lan et al. [50] digests the text of a mathematical equation provided as the answer to a free-response question. It transforms the text of an equation into a vector space of features, allowing these free responses to be clustered. Depending on the cluster that an equation is a part of, the equation can either be marked as correct, or MLP can identify the specific error made in the equation so that detailed feedback can be provided and partial credit can be awarded.

Automated systems can also help evaluate students more frequently, and in more piecemeal intervals, than might be feasibly done by a human instructor. Basu et al. built OK, a system that asks students about how their submitted computer programs should behave before the program is evaluated for correctness [1]. This intervention aligns the student's mental model of how the computer program is implemented with how a correct version of the program should work. OK reduced student questions about one assignment by 79%, suggesting that rubrics and specifications may not be enough to completely articulate an assignment's requirements.

Finally, several systems aim to predict the best way to guide students through different learning paths. Pardos et al. modified the edX platform to log user data and to suggest page destinations based on student behavior [64]. Piech et al. used student data from Code.org and Markov models to develop strategies to prevent students from getting stuck during a course [65]. Furthermore, O'Rourke et al. present evidence that intervening with hints may actually detract from student learning [59]. The question of when and how to intervene during a learning experience is still an area open for exploration.

Reflection: Embedding intelligence into a system can allow students to experiment with course concepts. This creates a safer environment for failure, so that their errors can be corrected quickly, privately, and without serious consequences.

INTERACTIONS WITH SOCIOTECHNICAL SYSTEMS

A great deal of technical software infrastructure underlies most online learning environments, but the social infrastructure that emerges when scaling learning to global populations has also been a major research focus. The interplay between these forces creates a sociotechnical system with emergent properties. The next set of Learning at Scale papers from Table 1 present empirical studies of learner interactions with such systems. In this section we survey a sample of representative papers related to both individual and group interactions.

Individual Interactions with Learning Technologies

One major research direction here has focused on understanding how students behave with respect to their use of the resources and digital artifacts present in many online courses. There have been several attempts at characterizing the general patterns of student interactions with these technologies. Wilkowski et al. present four simple categories to partition MOOC student personas: no-shows, observers, casual learners, and completers [84]. Though much of popular press coverage of online courses has been concerned with their low completion rates [62], Wilkowski et al. re-frame the goal of an online course as being meant for providing relevant learning opportunities at whatever levels students are willing to engage.

One particular behavior derived by Wilkowski et al. from data they collected is that students will complete assignments without looking at course materials until they reach an assignment they cannot complete. A similar observation was made by Guo and Reinecke when studying MOOC navigation behaviors [29]. Even students who would eventually earn a certificate skipped 22% of course content. One of the most common behaviors among certificate earners was navigating from an assessment back to a lecture, suggesting that they did not fully understand lecture material before they looked at the assessment, or they skipped the lecture altogether. The authors surmise that this behavior may reveal students' motivations to earn a credential rather than to truly engage with the material.

Since videos comprised the bulk of MOOC lecture content [28], they are often studied as conduits of behavioral analysis. For instance, Kim et al. performed a detailed study of students' navigation through lecture videos [36]. They examine peaks of activity within videos as a mechanism for identifying how students use video content. Over 50% of all videos are not watched until the end, and long videos have higher dropout rates. Kim et al. also identified several different student actions within videos including returning to missed content, following a tutorial step, and repeating a non-visual explanation. This work raises the open question of whether video content and playback can be optimized for different student goals.

Reflection: Students come into online courses with diverse goals, so course designers should make materials easy to access along non-linear and incomplete paths.

Group Interactions within Online Communities

The communities that form during online learning experiences involve a spectrum of participants from novices to experts, sometimes with the explicit goal of providing mentorship opportunities.

Researchers have observed that students want to contribute to a course even after they have completed it. Joyner saw that online students feel a sense of ownership of the program that they participate in, and this feeling motivates students to become teaching assistants [34]. Nelimarkka et al. conducted interviews with "alumni" MOOC participants. They found that alumni helped new students, and that these alumni were motivated by the desire to learn more from other students [56].

Zheng et al. examined the social dynamics of a MOOC, which they model as a "virtual organization" [89]. They collected data from Facebook groups organized around several Coursera courses, finding that forum post engagement on Facebook was higher compared to the official Coursera forums for the same course. Through a series of interviews with students, Zheng et al. found that students felt that it was easier to make new friends on Facebook compared to Coursera, and that there was a greater sense of trust between students on Facebook since accounts were connected to their real identities. Students reported that they felt their interactions were more "real" on Facebook, and that Facebook provided a platform that was easier for organizing group work compared to Coursera. Zheng et al.'s study suggests that serious improvements can be made to foster authenticity, trust, and other requirements for effective collaboration, sharing, and teaching in online classrooms.

Examining traditional discussion channels is a ripe starting point for examining the space of interactions in online learning communities. Coetzee et al. studied chatrooms in MOOCs, and although students report that the presence of a chatroom helps them in the course, there is little evidence that chatrooms provide any measurable performance benefits to students [10]. Coetzee et al. deployed both an embedded chat window in a single-page version of the course, and a chat tab in another version. They found that the embedded chat was used significantly more often. Through interviews they learned that many students liked to use the chat right before an assignment was due, and although a minority of students used the chat, many students reported that reading the discussions between other students was helpful. Students posted hundreds of links in the chatroom, including links to the course's own forum and links to code samples. This study opens up questions of how the creation and curation of knowledge by students can be better harnessed for wider community benefit.

Lastly, online communities of learning have subtleties of structure that may affect how students engage with the community. Kotturi et al. explored these ideas by deploying different social and learning technologies in a set of online courses [41]. They found that merely making technologies available for facilitating online discussions and connections is not enough. It is important to establish norms of behavior. Norms were established by showing students exemplar samples of student feedback, and by contacting the students to remind them that their social participation in the course is important. Both of these strategies led to higher student engagement in the studied courses. This investigation highlights how technologies that directly connect students can help students understand that their individual presence and contribution matter, despite the isolating effects of distance learning technology.

Reflection: Learning is an inherently social experience and technologies can facilitate student-to-student interaction. But these technologies may be secondary to establishing a culture that encourages frequent and open communication.

UNDERSTANDING ONLINE STUDENTS

Learning *at scale* implies an activity that is distributed among a diverse worldwide population. Therefore, understanding the goals, needs, and obstacles of students from a wide variety of backgrounds has been an active field of study at this conference. In this synthesis of the final set of papers from Table 1, we discuss how student knowledge can be modeled in context, and how global access and equity can be promoted.

Modeling Student Knowledge

To paraphrase Senator Howard Baker, designers of online learning experiences are often interested in what students know, and when they know it [19]. Modeling the diverse array of student knowledge before, during, and after a course can help instructors understand how course material can be improved, when to intervene if a student is struggling, and whether or not a student is prepared for more learning.

For example, Chen et al. investigated whether or not students applied knowledge they gained in a MOOC after the conclusion of a course [9]. They tracked the public activity of GitHub accounts that had been used by students in a course on functional programming. After the course was over, they found that past students wrote functional-style code with more frequency than before the course. Their study demonstrated that following students' public activity after a course ends can open up new methods of longitudinally evaluating the effectiveness of online learning strategies.

The medium of the online classroom also allows the kind of action-tracking that is commonly associated with e-commerce websites; such an approach can be constructively adapted to improve the understanding of student learning paths. Fratamico et al. deployed a virtual electrical engineering environment in a course [23]. By logging every action that a student performed, they were able to visualize sequences of actions over time, including how student decisions resulted in divergent and convergent paths of action. The data from these naturally occurring paths can be used to inform future decisions about student knowledge scaffolding and course structure.

In a philosophically similar approach, Coleman et al. adapted the Latent Dirichlet Allocation (LDA) technique to characterize sets of student actions in a course [11]. Their approach accurately identifies latent patterns of student interaction with the course. These patterns can be visualized by amounts and types of student activity over time. This study demonstrates a method for understanding patterns of student behavior across the many different artifacts within a course, which can provide information about how students are understanding course content and where they are experiencing success and failure.

Reflection: Now that we have abundant techniques for modeling student knowledge, what are effective and practical methods to deploy them at scale and measure their efficacy?

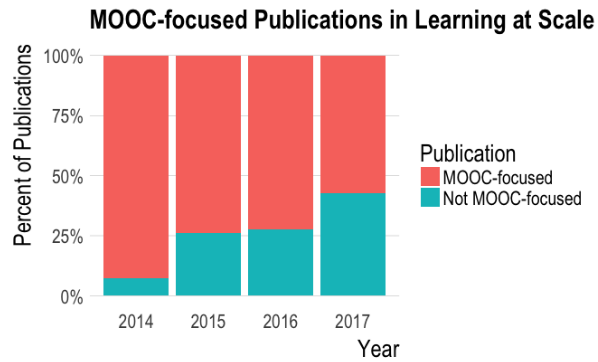

Figure 1. Learning at Scale has become less focused on MOOCs as the years progressed. MOOCs will still probably be an important part of online education in the foreseeable future, but we are just beginning to explore the variety of substrates that make scaled learning possible.

Promoting Global Access and Equity

Many projects to scale learning inherently have a social mission. The Learning at Scale community exists in a unique moment in history where we can imagine and oftentimes create educational experiences that have the potential to reach millions of people. Considering our privileged position and the global impact that this can have, it is important to integrate the values of equity, access, and openness into our work.

68 out of 69 Learning at Scale papers so far are authored by researchers working in Western (e.g., U.S. and Western European) institutions; the sole exception was from Microsoft Research India [13], which is still a U.S. company. It is important for researchers to consider that the vast majority of online students do not come from these same backgrounds [12, 21, 79]. Toward this end, Dasgupta et al. studied the effect of language localization on the speed at which students learned Scratch, a novice-oriented programming language [15]. With data from five localized Scratch learning communities, their analysis shows a small but significant increase in the speed at which students learned Scratch when they were learning in their native language. Making it easy for a learning experience to be localized (perhaps even by members of the community) should always be an important design consideration.

Furthermore, it is likely that the epoch of the computer desktop will look like a mere blip in time when the future history of computing is written, relative to the rapid spread and democratization of mobile devices. Notably, the only publication in Learning at Scale so far to specifically focus on using mobile devices for learning was a study by Breazeal et al. [3]. They developed and deployed a mobile application for promoting child literacy. They tested their system in Ethiopia, South Africa, and the United States in settings with a wide spectrum of access to technology. It took only a few days for all groups of children to quickly learn how to use the tablets provided to them. Their experiment pushed up against some limits of scale, though, since the devices had to be delivered by the research team, and their students were children with low literacy.

Scaling to a global audience also means reconciling drastically different worldviews amongst course participants. Conflict in online discussions is pervasive, but optimistically Reich et al. show that civil and productive discussion is possible in an online course [68]. Their education policy MOOC attracted students with diverse political beliefs, who participated in forums and individual threads with equal frequency, often endorsing posts by students with opposing political beliefs. Students with different beliefs also discussed the same controversial topics with the same level of frequency, suggesting that they were willing to engage with each other about issues. This course serves as a case study for practices and policies which foster empathic and insightful learning experiences.

Finally, even if students have the skills and resources to take part in an online course experience, there are still obstacles about their sense of personal identity that hold them back. Within the setting of a Chinese MOOC, Kizilcec et al. explored the role of Social Identity Threat (SIT): the fear that one will be negatively judged due to one's personal identity [37]. First they surveyed students about factors such as their gender, their perceived class in society, and their parents' socioeconomic status. Students were then randomized into being prompted to complete a writing assignment about either study skills (the control), or about their values and message they would send to their future selves (the intervention). Students with high SIT (e.g., lower social class men) who received the intervention completed the course at a significantly higher rate than those in the same group who did not receive the intervention.

Reflection: Modern learner populations are globally diverse, so it is important to empirically investigate new ideas about student identity and the settings in which they are learning.

DISCUSSION: CHARTING THE FUTURE

We conclude by reflecting on salient facets of our taxonomy and pointing out seven potentially fruitful directions for future research in the coming years of Learning at Scale.

1. Beyond the MOOC Hype Cycle

We believe that the wide array of insights published in Learning at Scale can be used to chart a trajectory for the growth and diversification of technology-enabled learning sciences. Back in 2011–2012, the three currently-largest MOOC providers (Coursera, edX, and Udacity) were founded [53, 16, 20],

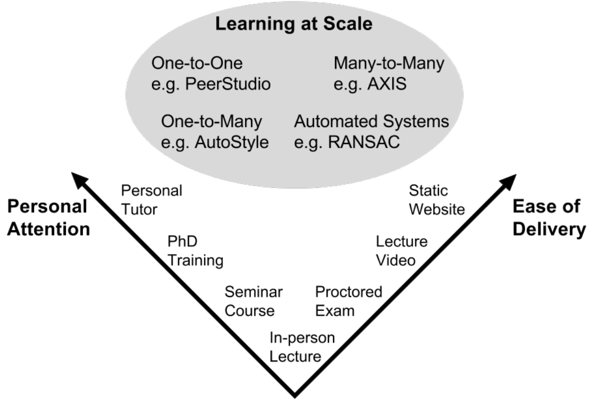

Figure 2. Personal attention and ease of delivery are two important considerations for designing educational experiences. One main goal of Learning at Scale is to contribute novel research that push toward the top of this graph with systems that maximize both dimensions of scale.

prompting the New York Times to declare "The Year of the MOOC" [62]. Following the model of the Gartner Hype Cycle [24], 2012 encompassed the peak of enthusiasm, but by 2013 Udacity co-founder Sebastian Thrun was already calling his own company a "lousy product" [6]. In a 2015 interview Coursera co-founder Daphne Koller stated that MOOCs were "emerging from the 'trough of disillusionment.'" [58]. This shifting balance is evident in Figure 1, where we plot the decreasing proportion of Learning at Scale publications each year that are focused on MOOCs.

What enduring lessons are there from this first phase of the MOOC hype cycle that can sustain forward-looking research programs in the coming years? One starting point to consider is that the MOOC revolution provided a valuable medium in which challenges for scaling learning could be found, addressed, and refined. It is now clear, though, that the strategies used to scale MOOCs have valuable applications outside of MOOCs, and that there are methods for scale that have been developed outside of MOOCs but could be incorporated back into them. This duality can provide guidance for future work.

2. The Axes of Scale for Online Learning Systems

Two important variables to be considered when evaluating the potential scale of educational experiences are the amount of personal attention required in order for the experience to be fully articulated, and how easily the experience can be delivered. For example, the experience of progressing through a Ph.D. program requires the personal attention of one or several advisors who focus on a small number of individuals in order to train them in a specific area of expertise. Each Ph.D. experience is unlike any other, but the attention of advisors is extraordinarily limited by funding and the amount of time they can spend training each of their pupils.

Compare the high level of personalization necessary for Ph.D. training to an experience like studying for a national standardized test, where there is no consideration to customize the test taking experience toward any of the attributes of the person taking the test. In addition to these exams being blind to the test-taker, the testing experience itself requires significant physical infrastructure to distribute: test-takers must travel to a prescribed location and then be monitored by proctors while they take the exam. This experience is both not personalized and fairly difficult to deliver at scale. Finally, consider the modern experience of learning mathematics with access to a computer or a smartphone. Any search engine will bring a student volumes of books, numerous hours of lecture videos, and thousands of static websites with explanations, strategies, and problem sets. The experience of learning mathematics in this way is trivially easy to distribute as long as both student and teacher have an internet connection; however these videos, books, and webpages have no knowledge of the student's background, interests, or intentions.

These two dimensions—personal attention and ease of delivery—combine into what we call the *Axes of Scale*, illustrated in Figure 2. We observe that the systems presented in Learning at Scale enable learning experiences that are easily and widely delivered with high levels of individual attention. This observation informed the organization of the first section of our taxonomy: Systems for Learning at Scale. Many of these systems can be conceptualized as Euclidean translations of existing methods towards the upper center of the Axes of Scale. For example: AXIS more widely distributes an expert explanation that might be provided by a personal tutor [86], PeerStudio organizes feedback like a student might get during an in-person seminar [46], and AutoStyle takes what could be a static style guide and instead offers an interactive experience for students [83]. We believe that the framework of Euclidean translation on the Axes of Scale can be used as a tool for inspiring new ideas about how to scale learning.

3. Toward More Immersive Online Course Experiences

One goal of the taxonomy that we have developed in Table 1 is to characterize some of the necessary ingredients for building engaging scaled learning experiences. We believe that individual principles from each category can be extracted and re-combined in novel ways to seed more immersive learning experiences. For instance, our taxonomy identifies a cluster of insights about video content including how videos should be produced [28], how long they should last [36], and how students interact with videos containing embedded assessments [42]. Moving to another section of our taxonomy but continuing with the thread of video and assessment we can see from Koedinger et al. that students doing activities tend to be more successful than students only watching videos [39].

A polymerization of these somewhat contradictory ideas might lead to creating novel online learning experiences beyond what is seen in standard MOOCs. For example, the Executive Data Science Specialization Capstone, a Coursera course, puts students into the shoes of a new data science manager [51]. By combining video interactions with real data scientists and simulated emails, documents, and data products, this course allows students to role play in an immersive fashion.

4. The Future is in Mobile Devices

One necessity for effectively scaling education is making sure that students have the technology available to access courses and educational content. All Learning at Scale papers so far (except for one [3]) focus on education delivered via traditional computers, but the nature of computing devices, and how people spend time on those devices, is quickly changing. In the last six months of 2016 global mobile internet usage surpassed desktop internet use [75]. The adoption of smartphones is accelerating all over the world: In 2015 68% of adults owned a smartphone in advanced economies, compared to 37% in emerging economies [5]. The three major MOOC providers, Coursera, edX, and Udacity, all publish their own mobile applications for both iOS and Android devices [12, 21, 79]. Despite the rise of mobile computing in the past few years, only one paper in our taxonomy specifically evaluates a scaled learning approach targeted for mobile devices [3]. Unlike a desktop or laptop, 94% of American smartphone owners say they carry their phone with them "frequently" [66]. How can we design for learning at scale when every student has a classroom in their pocket? With growing mobile internet use it is necessary to follow students to the devices they are using in order to evaluate the challenges and opportunities of learning on a mobile device rather than assuming that they will always be learning at their desks.

5. Addressing More Diverse Student Personas

Another trend that our taxonomy reveals is that there is widespread interest in categorizing students by their behavioral traces [84, 32, 11] and simultaneously there is general interest in tracking student learning trajectories [65, 22, 17, 23, 64]. How should we design online learning experiences considering that we can characterize students by behaviors and goals? So far the approach taken has been to try to predict student behavior and then to intervene, perhaps with a suggestion, like in Pardos et al. [64]. However, we know from user personas in HCI and prior studies that students use course materials according to their personal goals, which often do not include completing an entire course. The existence of these personas serves as a rebuttal to the common criticism of low completion rates in MOOCs: not every student who is seeking learning opportunities is looking to complete a course. As designers of courses and learning systems we should embrace this fact, especially given the diverse needs of online learners. This fact begs the question: Are there scalable approaches for explicitly guiding students according to their persona? Mullaney et al. observe that many students only ever see the beginning of a course [55]. Therefore, perhaps students should be offered guidance at the beginning of a course about how to best use the materials depending on what their goals are. In addition to being given explicit direction, inferences from student behavioral data could be used to offer feedback and encouragement to students who aspire to complete a course. Allowing students to visualize their own path through a course could better help them plan for keeping up with course demands.

6. Students as Active Co-Creators of Courses

Students are by no means passive participants in online learning experiences. Students curate knowledge, including course-relevant web links as explored by Coetzee et al. in their analysis of chatroom activity [10]. Huang et al. report that students who post in MOOC forums with the highest volume tend to create posts that are high in quality and the presence of their posts are correlated with higher overall forum activity [32]. Interviews conducted by Nelimarkka et al. show that past course participants are altruistically motivated to stay in course chatrooms to help new students [56]. Students also act as testers of learning materials and experiences, since they report inaccuracies or inconsistencies if they can find them. In several ways open courses share the properties of open source software, insofar as they are large collaborative projects in which many people participate and contribute to varying degrees. A series of MOOCs on Coursera called the Data Science Specialization takes this open source metaphor even farther [44]. Students are encouraged to create their own tutorials and data analyses, which they can then contribute to a community website using GitHub. This mechanism allows the students to actively contribute back to the course community, and it gives them the opportunity to practice skills taught in the specialization. Clearly there is a demonstrated desire by students who want to contribute to courses, but systems for organizing these contributions have not yet materialized. Perhaps there is a way to reconcile the visions of early connectivist cMOOCs with the modern xMOOC platforms that are now pervasive [30].

7. Fostering Better Social Connections in Online Courses

One of the biggest and most concerning consequences of scaled education is that it is more difficult for students to develop meaningful relationships compared to in-person educational settings. Education is much more than just the delivery of information; traditionally it has been an intensely social experience with students and instructors sharing a physical space. Online learning platforms commonly have forums, but forums often fall short of students' social needs. Zheng et al. report that these forums do not facilitate collaboration, nor do they allow students to build trust with each other compared to communicating on a Facebook group [89]. Facebook groups are not specifically designed for group learning, suggesting there are unexplored possibilities for creating environments that better serve groups of learners. To this end, Kulkarni et al. built Talkabout, a system for scaling and structuring small group video discussions [45]. They found that students who participated in these discussions were more engaged with the course and performed better. Kotturi et al. further tested Talkabout, noting that the mere option to use the system was not enough to motivate students to participate [41]. Expectations needed to be set via regular communications with students so that using Talkabout became a normal component of class activity. Even if there are more social options available compared to forums, students can benefit from the establishment of course culture that encourages face-to-face interactions.

Acknowledgments: Thanks to Kandarp Khandwala, Xiong Zhang, and the UC San Diego Design Lab for their feedback.

REFERENCES

1. Soumya Basu, Albert Wu, Brian Hou, and John DeNero. 2015. Problems Before Solutions: Automated Problem Clarification at Scale *(L@S '15)*. 205–213.
2. Anthony Botelho, Hao Wan, and Neil Heffernan. 2015. The Prediction of Student First Response Using Prerequisite Skills *(L@S '15)*. 39–45.
3. Cynthia Breazeal, Robin Morris, Stephanie Gottwald, Tinsley Galyean, and Maryanne Wolf. 2016. Mobile Devices for Early Literacy Intervention and Research with Global Reach *(L@S '16)*. 11–20.

4. Michael Brooks, Sumit Basu, Charles Jacobs, and Lucy Vanderwende. 2014. Divide and Correct: Using Clusters to Grade Short Answers at Scale *(L@S '14)*. 89–98.
5. Pew Research Center. 2016. Smartphone Ownership and Internet Usage Continues to Climb in Emerging Economies. (2016).
6. Max Chafkin. 2013. Udacity's Sebastian Thrun, Godfather Of Free Online Education, Changes Course. Fast Company. (2013).
7. John Champaign, Kimberly F. Colvin, Alwina Liu, Colin Fredericks, Daniel Seaton, and David E. Pritchard. 2014. Correlating Skill and Improvement in 2 MOOCs with a Student's Time on Tasks *(L@S '14)*. 11–20.
8. Binglin Chen, Matthew West, and Craig Zilles. 2017. Do Performance Trends Suggest Wide-spread Collaborative Cheating on Asynchronous Exams? *(L@S '17)*. 111–120.
9. Guanliang Chen, Dan Davis, Claudia Hauff, and Geert-Jan Houben. 2016. Learning Transfer: Does It Take Place in MOOCs? An Investigation into the Uptake of Functional Programming in Practice *(L@S '16)*. 409–418.
10. Derrick Coetzee, Armando Fox, Marti A. Hearst, and Bjoern Hartmann. 2014. Chatrooms in MOOCs: All Talk and No Action *(L@S '14)*. 127–136.
11. Cody A. Coleman, Daniel T. Seaton, and Isaac Chuang. 2015. Probabilistic Use Cases: Discovering Behavioral Patterns for Predicting Certification *(L@S '15)*. 141–148.
12. Coursera. 2018. `https://www.coursera.org/about/mobile`. (2018). Accessed: 2018-01-17.
13. Edward Cutrell, Jacki O'Neill, Srinath Bala, B. Nitish, Andrew Cross, Nakull Gupta, Viraj Kumar, and William Thies. 2015. Blended Learning in Indian Colleges with Massively Empowered Classroom *(L@S '15)*. 47–56.
14. Steven Dang, Michael Yudelson, and Kenneth R. Koedinger. 2017. Detecting Diligence with Online Behaviors on Intelligent Tutoring Systems *(L@S '17)*. 51–59.
15. Sayamindu Dasgupta and Benjamin Mako Hill. 2017. Learning to Code in Localized Programming Languages *(L@S '17)*. 33–39.
16. Nick DeSantis. 2012. Stanford Professor Gives Up Teaching Position, Hopes to Reach 500,000 Students at Online Start-Up. *The Chronicle of Higher Education* (2012).
17. Shayan Doroudi, Vincent Aleven, and Emma Brunskill. 2017. Robust Evaluation Matrix: Towards a More Principled Offline Exploration of Instructional Policies *(L@S '17)*. 3–12.
18. Nicholas Dronen, Peter W. Foltz, and Kyle Habermehl. 2015. Effective Sampling for Large-scale Automated Writing Evaluation Systems *(L@S '15)*. 3–10.
19. Will Dunham. 2014. Former Senator Howard Baker, key figure in Watergate hearings, dies. Reuters. (2014).
20. edX. 2017. About Us. (2017). `https://www.edx.org/about-us`
21. edX. 2018. `https://www.edx.org/mobile`. (2018). Accessed: 2018-01-17.
22. Mohammad H. Falakmasir, José P. Gonzalez-Brenes, Geoffrey J. Gordon, and Kristen E. DiCerbo. 2016. A Data-Driven Approach for Inferring Student Proficiency from Game Activity Logs *(L@S '16)*. 341–349.
23. Lauren Fratamico, Sarah Perez, and Ido Roll. 2017. A Visual Approach Towards Knowledge Engineering and Understanding How Students Learn in Complex Environments *(L@S '17)*. 13–22.
24. Gartner. 2018. Gartner Hype Cycle. `https://www.gartner.com/technology/research/methodologies/hype-cycle.jsp`. (2018). Accessed: 2018-01-17.
25. Chase Geigle, ChengXiang Zhai, and Duncan C. Ferguson. 2016. An Exploration of Automated Grading of Complex Assignments *(L@S '16)*. 351–360.
26. Ben U. Gelman, Chris Beckley, Aditya Johri, Carlotta Domeniconi, and Seungwon Yang. 2016. Online Urbanism: Interest-based Subcultures As Drivers of Informal Learning in an Online Community *(L@S '16)*. 21–30.
27. Yue Gong and Joseph E. Beck. 2015. Towards Detecting Wheel-Spinning: Future Failure in Mastery Learning *(L@S '15)*. 67–74.
28. Philip J. Guo, Juho Kim, and Rob Rubin. 2014. How Video Production Affects Student Engagement: An Empirical Study of MOOC Videos *(L@S '14)*. 41–50.
29. Philip J. Guo and Katharina Reinecke. 2014. Demographic Differences in How Students Navigate Through MOOCs *(L@S '14)*. 21–30.
30. Jonathan Haber. 2013. xMOOC vs. cMOOC. `http://degreeoffreedom.org/xmooc-vs-cmooc/`. (2013).
31. Andrew Head, Elena Glassman, Gustavo Soares, Ryo Suzuki, Lucas Figueredo, Loris D'Antoni, and Björn Hartmann. 2017. Writing Reusable Code Feedback at Scale with Mixed-Initiative Program Synthesis *(L@S '17)*. 89–98.
32. Jonathan Huang, Anirban Dasgupta, Arpita Ghosh, Jane Manning, and Marc Sanders. 2014. Superposter Behavior in MOOC Forums *(L@S '14)*. 117–126.
33. Petr Johanes. 2017. Epistemic Cognition: A Promising and Necessary Construct for Enriching Large-scale Online Learning Analysis *(L@S '17)*. 61–69.
34. David A. Joyner. 2017. Scaling Expert Feedback: Two Case Studies *(L@S '17)*. 71–80.
35. David A. Joyner, Wade Ashby, Liam Irish, Yeeling Lam, Jacob Langston, Isabel Lupiani, Mike Lustig, Paige Pettoruto, Dana Sheahen, Angela Smiley, Amy Bruckman, and Ashok Goel. 2016. Graders As Meta-Reviewers: Simultaneously Scaling and Improving Expert Evaluation for Large Online Classrooms *(L@S '16)*. 399–408.
36. Juho Kim, Philip J. Guo, Daniel T. Seaton, Piotr Mitros, Krzysztof Z. Gajos, and Robert C. Miller. 2014. Understanding In-video Dropouts and Interaction Peaks Inonline Lecture Videos *(L@S '14)*. 31–40.
37. René F. Kizilcec, Glenn M. Davis, and Geoffrey L. Cohen. 2017. Towards Equal Opportunities in MOOCs: Affirmation Reduces Gender & Social-Class Achievement Gaps in China *(L@S '17)*. 121–130.
38. René F. Kizilcec and Sherif Halawa. 2015. Attrition and Achievement Gaps in Online Learning *(L@S '15)*. 57–66.
39. Kenneth R. Koedinger, Jihee Kim, Julianna Zhuxin Jia, Elizabeth A. McLaughlin, and Norman L. Bier. 2015. Learning is Not a Spectator Sport: Doing is Better Than Watching for Learning from a MOOC *(L@S '15)*. 111–120.
40. Joseph A. Konstan, J.D. Walker, D. Christopher Brooks, Keith Brown, and Michael D. Ekstrand. 2014. Teaching Recommender Systems at Large Scale: Evaluation and Lessons Learned from a Hybrid MOOC *(L@S '14)*. 61–70.
41. Yasmine Kotturi, Chinmay E. Kulkarni, Michael S. Bernstein, and Scott Klemmer. 2015. Structure and Messaging Techniques for Online Peer Learning Systems That Increase Stickiness *(L@S '15)*. 31–38.
42. Geza Kovacs. 2016. Effects of In-Video Quizzes on MOOC Lecture Viewing *(L@S '16)*. 31–40.
43. Markus Krause, Marc Mogalle, Henning Pohl, and Joseph Jay Williams. 2015. A Playful Game Changer: Fostering Student Retention in Online Education with Social Gamification *(L@S '15)*. 95–102.
44. Sean Kross, Roger D. Peng, Brian S. Caffo, Ira Gooding, and Jeffery T. Leek. 2017. The democratization of data science education. *PeerJ Preprints* (2017).
45. Chinmay Kulkarni, Julia Cambre, Yasmine Kotturi, Michael S. Bernstein, and Scott R. Klemmer. 2015b. Talkabout: Making Distance Matter with Small Groups in Massive Classes *(CSCW '15)*. 1116–1128.
46. Chinmay E. Kulkarni, Michael S. Bernstein, and Scott R. Klemmer. 2015a. PeerStudio: Rapid Peer Feedback Emphasizes Revision and Improves Performance *(L@S '15)*. 75–84.
47. Chinmay E. Kulkarni, Richard Socher, Michael S. Bernstein, and Scott R. Klemmer. 2014. Scaling Short-answer Grading by Combining Peer Assessment with Algorithmic Scoring *(L@S '14)*. 99–108.

48. Youngwook Paul Kwon and Sara McMains. 2015. An Automated Grading/Feedback System for 3-View Engineering Drawings Using RANSAC *(L@S '15)*. 157–166.

49. Anne Lamb, Jascha Smilack, Andrew Ho, and Justin Reich. 2015. Addressing Common Analytic Challenges to Randomized Experiments in MOOCs: Attrition and Zero-Inflation *(L@S '15)*. 21–30.

50. Andrew S. Lan, Divyanshu Vats, Andrew E. Waters, and Richard G. Baraniuk. 2015. Mathematical Language Processing: Automatic Grading and Feedback for Open Response Mathematical Questions *(L@S '15)*. 167–176.

51. Jeffery T. Leek, Roger D. Peng, Brian S. Caffo, Ira Gooding, Sean Kross, and Jessica Crowell. 2016. Executive Data Science Capstone. **https://www.coursera.org/learn/executive-data-science-capstone**. (2016). Accessed: 2018-01-17.

52. Juho Leinonen, Petri Ihantola, and Arto Hellas. 2017. Preventing Keystroke Based Identification in Open Data Sets *(L@S '17)*. 101–109.

53. Tamar Lewin. 2012. Universities Reshaping Education on the Web. *The New York Times* (2012).

54. Harold A Linstone, Murray Turoff, and others. 1975. *The delphi method*. Addison-Wesley Reading, MA.

55. Tommy Mullaney and Justin Reich. 2015. Staggered Versus All-At-Once Content Release in Massive Open Online Courses: Evaluating a Natural Experiment *(L@S '15)*. 185–194.

56. Matti Nelimarkka and Arto Vihavainen. 2015. Alumni & Tenured Participants in MOOCs: Analysis of Two Years of MOOC Discussion Channel Activity *(L@S '15)*. 85–93.

57. Benjamin D Nye. 2015. Intelligent tutoring systems by and for the developing world: A review of trends and approaches for educational technology in a global context. *International Journal of Artificial Intelligence in Education* 25, 2 (2015), 177–203.

58. The University of Pennsylvania. 2015. The Hype is Dead, but MOOCs Are Marching On. **http://knowledge.wharton.upenn.edu/article/moocs-making-progress-hype-died/**. (2015). Accessed: 2018-01-17.

59. Eleanor O'Rourke, Christy Ballweber, and Zoran Popovií. 2014. Hint Systems May Negatively Impact Performance in Educational Games *(L@S '14)*. 51–60.

60. Eleanor O'Rourke, Erin Peach, Carol S. Dweck, and Zoran Popovic. 2016. Brain Points: A Deeper Look at a Growth Mindset Incentive Structure for an Educational Game *(L@S '16)*. 41–50.

61. Korinn Ostrow, Christopher Donnelly, Seth Adjei, and Neil Heffernan. 2015. Improving Student Modeling Through Partial Credit and Problem Difficulty *(L@S '15)*. 11–20.

62. Laura Pappano. 2012. The Year of the MOOC. *The New York Times* (2012).

63. Zachary A. Pardos and Kevin Kao. 2015. moocRP: An Open-source Analytics Platform *(L@S '15)*. 103–110.

64. Zachary A. Pardos, Steven Tang, Daniel Davis, and Christopher Vu Le. 2017. Enabling Real-Time Adaptivity in MOOCs with a Personalized Next-Step Recommendation Framework *(L@S '17)*. 23–32.

65. Chris Piech, Mehran Sahami, Jonathan Huang, and Leonidas Guibas. 2015. Autonomously Generating Hints by Inferring Problem Solving Policies *(L@S '15)*. 195–204.

66. Lee Rainie and Kathryn Zickuhr. 2015. Americans' Views on Mobile Etiquette. (2015).

67. Karthik Raman and Thorsten Joachims. 2015. Bayesian Ordinal Peer Grading *(L@S '15)*. 149–156.

68. Justin Reich, Brandon Stewart, Kimia Mavon, and Dustin Tingley. 2016. The Civic Mission of MOOCs: Measuring Engagement Across Political Differences in Forums *(L@S '16)*. 1–10.

69. Steve Ritter, Michael Yudelson, Stephen E. Fancsali, and Susan R. Berman. 2016. How Mastery Learning Works at Scale *(L@S '16)*. 71–79.

70. Jose A. Ruiperez-Valiente, Giora Alexandron, Zhongzhou Chen, and David E. Pritchard. 2016. Using Multiple Accounts for Harvesting Solutions in MOOCs *(L@S '16)*. 63–70.

71. Mehdi S.M. Sajjadi, Morteza Alamgir, and Ulrike von Luxburg. 2016. Peer Grading in a Course on Algorithms and Data Structures: Machine Learning Algorithms Do Not Improve over Simple Baselines *(L@S '16)*. 369–378.

72. Arjun Singh, Sergey Karayev, Kevin Gutowski, and Pieter Abbeel. 2017. Gradescope: A Fast, Flexible, and Fair System for Scalable Assessment of Handwritten Work *(L@S '17)*. 81–88.

73. Sumukh Sridhara, Brian Hou, Jeffrey Lu, and John DeNero. 2016. Fuzz Testing Projects in Massive Courses *(L@S '16)*. 361–367.

74. Adam Stankiewicz and Chinmay Kulkarni. 2016. $1 Conversational Turn Detector: Measuring How Video Conversations Affect Student Learning in Online Classes *(L@S '16)*. 81–88.

75. StatCounter. 2016. Mobile and tablet internet usage exceeds desktop for first time worldwide. **http://gs.statcounter.com/press/mobile-and-tablet-internet-usage-exceeds-desktop-for-first-time-worldwide**. (2016). Accessed: 2018-01-17.

76. Thomas Staubitz, Dominic Petrick, Matthias Bauer, Jan Renz, and Christoph Meinel. 2016. Improving the Peer Assessment Experience on MOOC Platforms *(L@S '16)*. 389–398.

77. Kristin Stephens-Martinez, Marti A. Hearst, and Armando Fox. 2014. Monitoring MOOCs: Which Information Sources Do Instructors Value? *(L@S '14)*. 79–88.

78. Jonathan H. Tomkin and Donna Charlevoix. 2014. Do Professors Matter?: Using an a/B Test to Evaluate the Impact of Instructor Involvement on MOOC Student Outcomes *(L@S '14)*. 71–78.

79. Udacity. 2018. **https://www.udacity.com/mobile**. (2018). Accessed: 2018-01-17.

80. Frans Van der Sluis, Jasper Ginn, and Tim Van der Zee. 2016. Explaining Student Behavior at Scale: The Influence of Video Complexity on Student Dwelling Time *(L@S '16)*. 51–60.

81. Frans van der Sluis, Tim van der Zee, and Jasper Ginn. 2017. Learning About Learning at Scale: Methodological Challenges and Recommendations *(L@S '17)*. 131–140.

82. Andrew E. Waters, David Tinapple, and Richard G. Baraniuk. 2015. BayesRank: A Bayesian Approach to Ranked Peer Grading *(L@S '15)*. 177–183.

83. Eliane S. Wiese, Michael Yen, Antares Chen, Lucas A. Santos, and Armando Fox. 2017. Teaching Students to Recognize and Implement Good Coding Style *(L@S '17)*. 41–50.

84. Julia Wilkowski, Amit Deutsch, and Daniel M. Russell. 2014a. Student Skill and Goal Achievement in the Mapping with Google MOOC *(L@S '14)*. 3–10.

85. Julia Wilkowski, Daniel M. Russell, and Amit Deutsch. 2014b. Self-evaluation in Advanced Power Searching and Mapping with Google MOOCs *(L@S '14)*. 109–116.

86. Joseph Jay Williams, Juho Kim, Anna Rafferty, Samuel Maldonado, Krzysztof Z. Gajos, Walter S. Lasecki, and Neil Heffernan. 2016. AXIS: Generating Explanations at Scale with Learnersourcing and Machine Learning *(L@S '16)*. 379–388.

87. David R. Woolley. 1994. **http://www.thinkofit.com/plato/dwplato.htm**. (1994). Accessed: 2018-01-17.

88. Diyi Yang, Miaomiao Wen, Iris Howley, Robert Kraut, and Carolyn Rose. 2015. Exploring the Effect of Confusion in Discussion Forums of Massive Open Online Courses *(L@S '15)*. 121–130.

89. Saijing Zheng, Kyungsik Han, Mary Beth Rosson, and John M. Carroll. 2016. The Role of Social Media in MOOCs: How to Use Social Media to Enhance Student Retention *(L@S '16)*. 419–428.