Chapter 18: Learning Analytics and Learning at Scale

Justin Reich[1]

[1] Teaching Systems Lab, Massachusetts Institute of Technology, Cambridge, USA

DOI: 10.18608/hla22.018

ABSTRACT

Learning at scale – an interdisciplinary field at the intersection of learning science and computer science – investigates learning environments with many, many learners and few experts to guide them. In recent decades, new large-scale learning environments have been announced with much fanfare about their potential to transform or "disrupt" traditional systems of formal schooling. This disruption has not occurred. Rather, new technologies are put to use in limited ways in specific niches of the existing education system, and the growth in their adoption is more steady and linear than abrupt or exponential. Though the societal impact of learning at scale has been uneven and incremental, the best hope for making the most of new large-scale technologies is through a continuous process of research and improvement. (The ideas in this chapter are expanded on in [24]).

Keywords: Learning at scale, learning analytics, peer learning, adaptive tutors, massive open online courses, design-based research, experimental research

INTRODUCTION

Many hoped that the massive volumes of fine-grained, global-scale learning tracking data, when combined with new forms of computational analysis, would lead to data-driven breakthroughs in learning science or instructional design (See for instance, [15]). This dream has not come to fruition. To date, learning at scale research has led to some useful insights on what might be called "educational policy analytics" – studies of how learners from different life circumstances use learning technologies differently – and "education behavior analytics"—how people click and act in online learning platforms. But research insights about learning—about changes in human cognition or capacity – from studies of large-scale technologies have been far more limited. The most promising possible future for learning analytics in learning at scale will not come from accumulating larger or more fine-grained troves of user data, but from research studies that use design-based or experimental methods to study systematic variation in competing approaches to effective design of large-scale learning [26].

In what follows, I address four questions, 1) What is learning at scale? 2) How has learning at scale changed the nature of education? 3) What has learning analytics and related research revealed about learning at scale? And 4) What are the possible futures for design and research in learning at scale?

1 WHAT IS LEARNING AT SCALE?

The ACM Learning@Scale 2020 conference home page offers a useful summary of the field:

> L@S investigates large-scale, technology-mediated learning environments that typically have many active learners and few experts on hand to guide their progress or respond to individual needs. Modern learning at scale typically draws on data at scale, collected from current learners and previous cohorts of learners over time. Large-scale learning environments are very diverse. Formal institutional education in K-16 and campus-based courses in popular fields involve many learners, relative to the number of teaching staff, and leverage varying forms of data collection and automated support. Evolving forms of massive open online courses, mobile learning applications, intelligent tutoring systems, open courseware, learning games, citizen science communities, collaborative programming communities (e.g. Scratch), community tutorial systems (e.g. StackOverflow), shared critique communities (e.g. DeviantArt), and countless informal communities of learners (e.g. the Explain It Like I'm Five sub-Reddit) are all examples of learning at scale. All share a common purpose to increase human potential, leveraging data collection, data analysis, human

interaction, and varying forms of computational assessment, adaptation and guidance.

The diverse learning environments described above can be categorized into three genres defined by the question, "Who sets the sequence of learning activities?" These sequences can be created by instructors – as in the case of MOOCs, by algorithms – as in the case of adaptive tutoring software, or by peers – as in the case of distributed learning networks. Each of these genres of instructor-guided, algorithm-guided, and peer-guided large-scale learning technologies has a history, a research literature, and a track record of success and failures in formal educational institutions. Each genre also uses a common set of core technologies, and they reenact pedagogical debates that have deep roots in the history of education. Figure 1 summarizes the three genres, and then I discuss the genres, their technologies, and their pedagogical roots below.

The massive open online courses (MOOCS) created by elite universities are examples of instructor-driven learning experiences [10]. Instructors design or select lectures, readings, and activities that form a knowledge base for student learning. Learners are assessed by tools and systems designed by instructors, that can range from simple multiple-choice questions to complex systems for evaluating computer programming assignments. The learning experiences in the course are arranged in a particular order, from the Shang Dynasty to the Era of Mao or from "Hello World" to recursive algorithms, that are selected by the instructor. A student may be free to traverse this material in her own way, and she might help a peer along the path, most students generally proceed along the main path laid out by instructors.

Adaptive, large-scale learning environments are those where each item in a learning sequence is selected by an algorithm or other system on the basis of student performance in previous parts of a learning sequence. These kinds of learning experiences are often called adaptive tutors or computer-assisted instruction, and Khan Academy offers a useful example. While Khan Academy is best known for Khan's video lectures, when Khan Academy is used in schools, students spend 85% of their time doing practice problems [21]. These problems will be familiar to anyone who has ever completed a worksheet in mathematics class. They pose a question, and students have to provide a correct answer by inputting an equation, selecting a point on a Cartesian plane, ordering a series of numbers in line, selecting from a list of multiple-choice options, and so forth. The problems are organized into topics, such as dividing fractions or solve quadratic equations.

Unlike a paper worksheet however, the order of problems that a student encounters depends upon her performance on each problem. Within a class of problems – such as multiplying fractions – some problems are easier (multiplying by ½) and some are harder (multiplying by 1/13). Students are given an initial problem, and if the student gets a problem right, an algorithm assigns a more difficult problem. If she gets it wrong, the system assigns an easier problem, perhaps along with some form of remediation, like a hint or link to an explanation. These systems are often called adaptive, since they can increase or decrease in difficulty and provide specific remediation based on the performance of the student. In nearly all MOOCs from edX or Coursera, every student receives the same number of problems and assignments which are presented in the same order. Students using Khan Academy and other adaptive tutors are offered a set of assignments that are dynamically adjusted for the individual student.

Peer-driven learning environments – like those proposed by Sugata Mitra in the School in the Cloud – are where participants can offer instruction, examples, comments and feedback, and users can follow each other, and form subgroups and networks. Mitra argued that if learners were organized into small groups, with access to the learning resources of the internet and some minimal on-demand mentoring and coaching (he proposed using a network of British pensioners in his trials), then students could learn any topic of any complexity [2, 20]. The original Connectivist-inspired MOOCs, provide another example of a peer-guided large-scale learning community. Participants created their own blogs, social media accounts, and other sites on the open web where they responded to course prompts and to each other. Instructors used the course home page and other technologies to aggregate copies of these diverse contributions into one central location, but at their most successful, peer interactions were the driving force of cMOOCs [19].

The most prominent peer-driven learning environment in K-12 schools is the community organized around the Scratch programming language, developed by the Lifelong Kindergarten Lab at MIT [29]. Scratch is a block-based programming language where the young and young-at-heart can learn to program by dragging "blocks" with executable code instructions into place with other blocks, rather than by writing programming syntax with specifications for spacing, semi-colons, variable names and so forth. By default, all Scratch programs exist as projects, all projects are publicly viewable and openly-licensed, and all projects can be forked and remixed as new projects, so that sharing and community are integral parts of the experience of using the Scratch programming language. In these communities there are designers and leaders—Mitch Resnick, Natalie Rusk, and many others in the Lifelong Kindergarten Lab create the environment for Scratchers to work and learn, they highlight projects on the Scratch website and social media, and cultivate community. This community then creates a wide array of projects, tutorials, guides, and other sub-communities, and learners in the Scratch community then choose for themselves how they navigate this web of opportunities for practice and learning.

The three genres of learning at scale – instructor, algorithm, and peer-guided – typically draw on different technologies, different pedagogies, and different research traditions. Instructor- and algorithm-guided large-scale learning environments typically depend upon some form of autograder to evaluate learner performance; by contrast

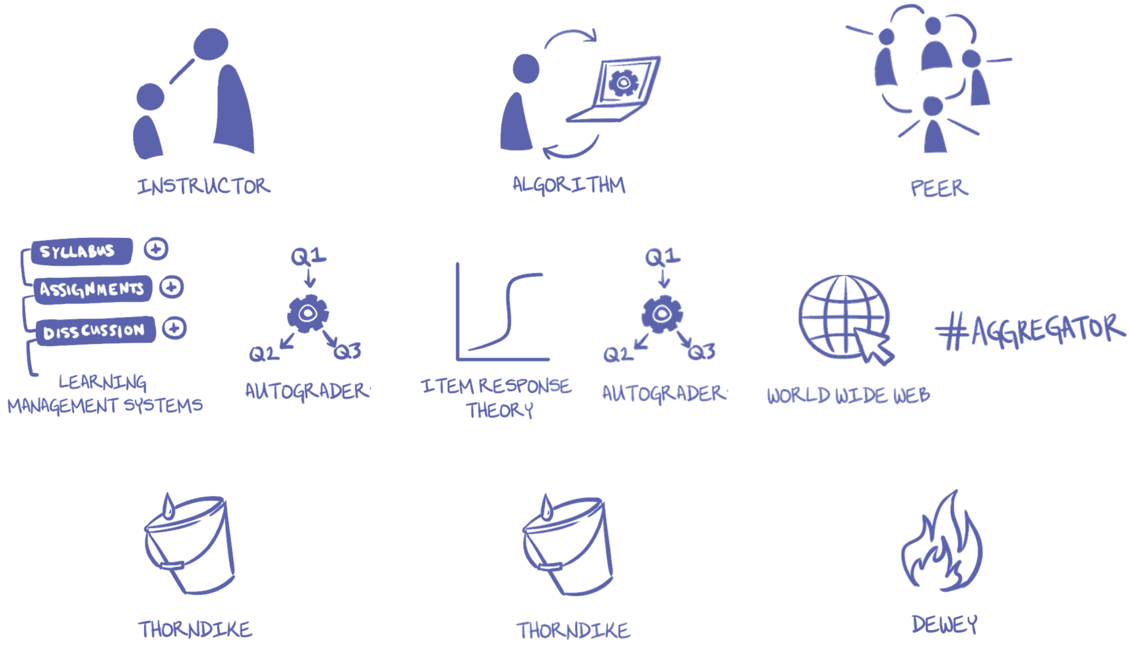

Figure 1: Genres of learning at scale.

peer-guided learning environments typically eschew formal assessment and focus on discourse and peer feedback. Instructor- and algorithm-guided genres typically take pedagogical inspiration from instructionist approaches to pedagogy, in the tradition of Thorndike [34] or more recently [33] where experts disseminate knowledge to be absorbed by novices. In the peer-guided genre, design is more often inspired by pedagogical philosophies emphasizing learner discovery and apprenticeship, like the Constructionism [6] at the heart of the Scratch programming community or the Connectivist [32] ideas that inspired the earliest massive open online courses. The three genres are also often studied by different research communities: scholars interested in adaptive tutors attend the International Conference on Artificial Intelligence and Adaptive Education or Educational Data Mining conference; those interested in instructor-guided learning at scale attend eMOOCs or Learning with MOOCs; and researchers studying peer-guided learning communities attend the Connected Learning Summit or the Constructionism conference.

Despite these differences, the three genres share much in common: they face a similar set of challenges in adoption in formal learning environments, and a common underlying data structure to track the activities of learners.

1.1 How has learning at scale changed the nature of education?

For those with access to global online networks, it is the greatest time in world history to be a learner. Never before have learners had such incredible access to resources, courses and communities of tutors and apprentices. Whether you want to learn to play guitar, brew beer, identify birds, translate Cicero, throw a javelin, intubate a trauma victim, integrate a function, detonate a bomb, program in Javascript, or become a better teacher, there are online classes, tutorials, forums, and networks full of people who are excited to teach and excited to learn. If you've ever signed up for an online class, downloaded an educational app, or watched a video about how to unclog a toilet, you are part of that network.

Yet, despite the extraordinary growth of informal online learning, changes to formal educational systems remain modest and targeted. Over the last twenty years, education technology advocates have promised dramatic changes in education systems. In 2008, Harvard Business School professor Clayton Christensen, with colleagues Curtis Johnson and Michael Horn [4], wrote a book called Disrupting Class about online learning and the future of K-12 schools. They predicted that in ten years – by 2019 – half of all middle and high school courses would be replaced by adaptive, self-paced online courses, and "the cost will be one-third of today's costs, and the courses will be much better." Udacity founder Sebastian Thrun argued that in 50 years, "there will be only 10 institutions in the world delivering higher education and Udacity has a shot at being one of them" [17]. Sugata Mitra went further to argue that in an internet-connected world, schools weren't even necessary:

> "Thirteen years of experiments in children's education takes us through a series of startling results – children can self organise their own learning, they can achieve educational objectives on their own, can read by themselves. Finally, the

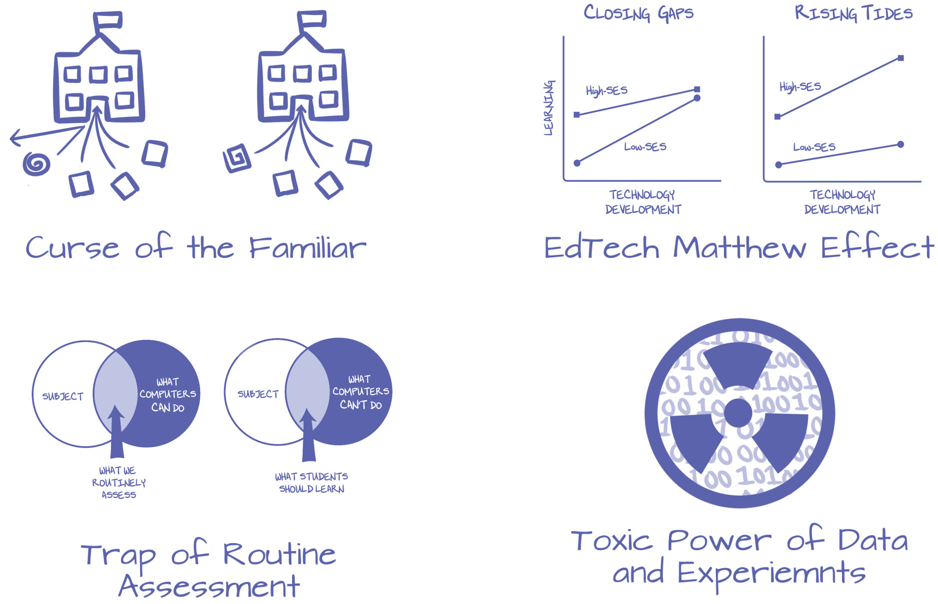

Figure 2: Four "As-Yet Intractable Dilemmas in Learning at Scale.

> most startling of them all: Groups of children with access to the Internet can learn anything by themselves" [20]

None of these predictions have come true, nor will they. The core misconception behind these predictions is that new technologies can disrupt, transform, or brush aside existing educational systems. This rarely, perhaps never, happens. Far more commonly, our complex, conservative educational systems domesticate new technologies, embedding them in existing routines in specific niches of the ecology of education.

One challenge to educational transformation is the "Curse of the Familiar" [23]. Educational systems can only readily adopt technologies that extend existing school practices. One of the most widely used educational websites in the world is Quizlet, which provides digital flashcards [8]. Tens of millions American students use Quizlet every year, but digitizing flashcards doesn't change routines in schools. Things which digitize existing practices can be readily adopted, but they provoke minimal changes in learning routines. By contrast, things which propose dramatic changes in learning routines are difficult to adopt. Early forms of Connectivist MOOCs offered a striking reinterpretation of learning practices in higher education, but many learners and instructors found their distributed, networked approaches to learning to be confusing [16].

Moreover, new technologies are typically only useful in specific subjects or disciplines. Both instructor-guided and algorithm-guided learning at scale technologies depend on autograders to computationally assess learner performance. Autograding technology, however is limited by what I call the "Trap of Routine Assessment." Computers are good at assessing the kinds of routine tasks that computers are good at doing, that we no longer need humans to do in the work force [27]. Autograders are good at assessing things with one right answer, or when a correct answer can be strictly defined by a set of decision rules. These are also the kinds of routine tasks that computers and robots can be programmed to accomplish. In math we have good autograders for computation, but not for explaining the reasoning behind computation strategies. In computer science, we have good autograders. In language arts, we have good autograders for the basics of decoding and pronunciation, but not for evaluating intrepretations of literature or poetry. The unevenness of our autograding technologies explains why large-scale learning technologies are more commonly found in some fields – STEM, computer science, early language acquisition – and not in others.

Like other education technologies, large-scale learning technologies typically disproportionately benefit the affluent. The "EdTech Matthew Effect" argues that like many sociological phenomenon, new technologies often accrue advantages to the already-advantaged [35, 28]. Morgan Ames [1] studied the roll out of One Laptop Per Child devices in Paraguay, and found that students who most deeply immersed themselves in the learning opportunities afforded by Scratch or Turtle Writer were those who had parents and families that had already introduced their children to learning opportunities with computers. MOOC researchers have consistently found that instructor-guided, large-scale learning depends on a well-developed set of self-regulated learning skills [11]. Already-affluent, already-educated learners are most likely to have had the opportunities to develop these skills, so MOOCs have not democratized education, but rather have accrued the bulk of their advantages to those who already had educational opportunities.

These common challenges help explain why the predictions of disruption and transformation from learning at scale have generally fallen flat. School systems are complex, technologies are uneven, opportunities are distributed inequitably in a highly-stratified society. Instead of dramatic transformations, we see specific technologies in specific disciplines used to the benefit of particular groups of users. If you are hoping that new technologies will be able to radically accelerate human development, the conclusion that change happens incrementally is probably a disappointment. But if you think that global human development is a game of inches – a slow, complex, maddening, plodding process with two steps back for every three steps forward – then the field of learning at scale offers one avenue for taking some of those forward steps.

1.2 What has learning analytics and related research revealed about learning at scale?

Across instructor-, algorithm-, and peer-guided learning environments, one of the unifying features of large-scale learning environments are the data and data structures that underlie these systems. At any given moment, a large-scale learning system needs to have a model of all possible actions that a learner can take – a model of the system – and a model of a student's state within this system. In Scratch, this might be all of the blocks assembled into a Scratcher's program at this particular moment; in a MOOC, this might mean tracking every assignment a student has completed to date and every assignment that is currently available but not yet completed. All of this data can be harnessed to create a complete record of what every learner has ever done within the system: a longitudinal record collected keystroke by keystroke and click by click, for millions of learners around the world. Large-scale learning environments are generating datasets that are orders of magnitude larger than what educational researchers have traditionally studied.

Coursera founder Daphne Koller [15] argued that these new sources would "turn the study of human learning from the hypothesis-driven mode to the data-driven mode, a transformation that, for example, has revolutionized biology." Since the founding of MOOCs, hundreds of millions of dollars have been spent on new courses, new platforms, and research efforts lead by some of the world's most accomplished computer scientists and learning scientists. Despite these efforts, Koller's prediction has not come to pass.

Researchers studying the vast new datasets from MOOCs have uncovered some useful findings about the demographics and behaviors of MOOC participants. For instance, despite an early rhetoric claiming that MOOCs could "democratize education," a number of studies have shown that people from more affluent countries and neighborhoods are more likely to register for MOOCs and once enrolled, more likely to complete them [9, 12]. Alongside these kinds of "educational policy analytics," much of the early research in MOOCs focused on correlations among behavioral measures. Deboer, Ho, Stump and Breslow [7] showed that a wide variety of learner inputs (videos viewed, problems answered, actions taken) correlated with each other and with outcomes like grade and earning a certification. Many studies published similar results, and I jokingly have described this line of inquiry as proving "Reich's Law," that students who do stuff do other stuff, and students who do stuff, do better than students who don't do stuff.

Two findings that go a step beyond Reich's law involve self-regulated learning, and the "doer" effect. Several MOOC studies found that successful learners showed evidence of proficiency with self-regulated learning, as measured by actions like reviewing prior material in the course [18, 11]. Given the very low levels of human supports available in MOOCs, these researchers theorize that proficiency with self-regulated learning is a prerequisite to success in MOOCs. Koedinger and colleagues [13, 14] at CMU showed in several studies that MOOC participants who engaged in problems and watched videos had better learning outcomes than students who only watched videos – a phenomenon they describe as the "doer effect." These are useful initial findings – that learners in courses without teachers need to be good students, and good students do problems and don't just watch videos—but they perhaps offer robust evidence for common sense, rather than new directions for the field of learning science. It turns out that researchers can collect terabytes of data about what people click without generating much new additional understanding of what's happening inside their heads.

Analytics researchers have also found it relatively straightforward to predict learner outcomes based on only a few initial weeks of user participation data [30, 38, 37]. Predicting who will drop out and succeed, however, is only useful to the extent that instructional designers can use that information to provide additional supports to struggling learners. To date, little research has shown how these predictions can be leveraged to improve student outcomes. Neil Heffernan, the principal investigator for the ASSISTments platform, an adaptive, math homework practice platform, once declared, "I now tell my students that no one is allowed to make a prediction without having some intervention planned to address the results of the prediction" [25]. Learning analytics without a linked intention to improve learning runs the risk of aimless fiddling.

1.3 What are the possible futures for research and design in learning at scale?

In his admonition to students, Heffernan anticipates one of the two sea changes in learning analytics research necessary for the field of learning at scale to advance. First, the case that learning science can be advanced by the passive, observational, cross-sectional study of massive datasets using advanced computational techniques thus far appears weak. Researchers need to be involved in designing studies that systematically introduce variation in instructional design to test the theory and practice of learning. In quantitative research traditions, this might look like randomized controlled trials that evaluate and compare differing instructional approaches. In qualitative research traditions, this might look more like iterative design-based research [31]. The massive, granular datasets collected by large-scale learning environments might prove especially useful in illuminating the mechanisms by which competing instructional designs might lead to better learning outcomes, but these large datasets need to be put in the service of design-based and experimental approaches, rather than more passive, observational, cross-sectional studies.

Second, the study of learning requires measures of learning. Most studies of large-scale learning platforms use measures and indicators derived from platform data, many of which are not well designed for tracking and evaluating learning. Studies of MOOCs use grades and certifications as proxies for learning, but many of these studies lack rigorous pre-test data (so it's not clear how much students are actually learning versus certifying preexisting competencies) and many of the assessments that under-gird these grades and certificates are not well designed. In peer-guided learning environments, the openended nature of learning environments provides another kind of assessment challenge – what does it mean to measure learning across Scratch projects if the point of Scratch is for young people to create whatever they want? Clever manipulation of the underlying activity data is no substitute for attention to these challenging issues of measurement. (Colvin and colleagues [5] offer one model of studying learning with well-validated measures in several physics courses).

Similarly, many studies of large-scale learning are bound entirely within a single platform, but one of the core purposes of learning is to transfer skills into new domains. Studying this transfer, therefore, is vital to understanding the potential and limits of learning at scale. A few studies have investigated transfer of learning "beyond the MOOC." To evaluate the impact of a Functional Programming MOOC, Chen, Davis, Hauff, and Houben [3] examined GitHub log data requests to find evidence of MOOC participants (using the same usernames across platforms) deploying programming skills from the MOOC in projects. To evaluate the impact of a course on learning analytics, Wang, Baker, and Pacquette [36] evaluated how MOOC participants joined scholarly societies and submitted papers in the field. Napier, Huttner-Loan, and Reich [22] studied how teachers adopted skills and practices from a MOOC about leading educational change. If one point of learning is to build human capacity to flexibly tackle future challenges, learning analytics will have to study students beyond learning platforms.

Contrary to predictions from the early days of MOOCs, the data collected by large-scale learning will not magically lead to a data-driven revolution in education science, but it still has potential to be a valuable resource in advancing learning science. The most promising future of learning analytics in large-scale learning will be interdisciplinary ventures conducted by joint teams of experts in substantive domains, in measurement and assessment, in design-based or experimental research, and in analyzing the granular data generated by large-scale platforms.

These efforts will not lead to the disruptive transformation of educational systems, but rather to steady, incremental progress in the field. Peer-guided learning technologies will be beloved platforms for devoted hobbyists – many, many children will get a brief introduction to computational creativity through Scratch, and a tiny handful will fall in love with the possibilities of the platform and blossom as programmers. Adaptive tutors will continue to find uses in educational systems in fields where human performance is amenable to evaluation by autograders, in fields like early language acquisition, mathematics, and computer science. Many students using adaptive tutors will learn a little more than they would have otherwise. MOOCs and other instructor-guided learning environments will primarily benefit those with the self-regulated learning skills to persevere through online learning with minimal supports; unfortunately most of the people who fall into this category are already-affluent, already-educated learners pursuing additional advanced credentials. In the status quo, large-scale learning is more likely to exacerbate educational inequality rather than to democratize education.

Learning analytics, learning at scale, and learning science as fields could all play a role in shifting this trajectory in a more positive, more equitable, and more promising direction. Such a shift would require embracing interdisciplinary research that recognizes the enormous complexity of iteratively improving systems that support learning at scale. It would require research that follows learners beyond online platforms and into the classrooms and workplaces where the transfer of skills can be observed and supported. It would require resisting the siren song of massive datasets and elegant, sophisticated post-hoc analysis, and reimagining large-scale learning analytics research in the service of more ambitious approaches to design-based and experimental research.

REFERENCES

[1] Morgan G. Ames. *The Charisma Machine*. The MIT Press, 2019. DOI: `10.7551/mitpress/10868.001.0001`. URL: `https://doi.org/10.7551%5C%2Fmitpress%2F10868.001.0001`.

[2] *Build a school in the cloud*. 2013. URL: `https://www.ted.com/talks/sugata_mitra_build_a_school_in_the_cloud`.

[3] Guanliang Chen, Dan Davis, Claudia Hauff, and Geert-Jan Houben. "Learning transfer: Does it take place in MOOCs? An investigation into the uptake of functional programming in practice". In: *Proceedings of the Third (2016) ACM Conference on Learning@ Scale*. 2016, pp. 409–418.

[4] Clayton M Christensen, Michael B Horn, and Curtis W Johnson. *Disrupting class: How disruptive in-*

novation will change the way the world learns. Vol. 1. McGraw-Hill New York, 2011.

[5] Kimberly F. Colvin, John Champaign, Alwina Liu, Qian Zhou, Colin Fredericks, and David E Pritchard. “Learning in an introductory physics MOOC: All cohorts learn equally, including an on-campus class”. In: *The International Review of Research in Open and Distributed Learning* 15.4 (May 2014). DOI: 10.19173/irrodl.v15i4.1902. URL: https://doi.org/10.19173%5C%2Firrodl.v15i4.1902.

[6] Peter G. Dean and Seymour Papert. “Mindstorms: Children, Computers and Powerful Ideas”. In: *The Mathematical Gazette* 65.434 (Dec. 1981), p. 298. DOI: 10.2307/3616611. URL: https://doi.org/10.2307%2F3616611.

[7] Jennifer DeBoer, Andrew D. Ho, Glenda S. Stump, and Lori Breslow. “Changing “Course””. In: *Educational Researcher* 43.2 (Mar. 2014), pp. 74–84. DOI: 10.3102/0013189x14523038. URL: https://doi.org/10.3102%5C%2F0013189x14523038.

[8] Matthew Glotzbach. *New monthly milestone for Quizlet: 50 million monthly learners*. 2018. URL: https://quizlet.com/blog/a-new-milestone-for-quizlet-50-million-monthly-%09%09learners.

[9] John D. Hansen and Justin Reich. “Democratizing education? Examining access and usage patterns in massive open online courses”. In: *Science* 350.6265 (Dec. 2015), pp. 1245–1248. DOI: 10.1126/science.aab3782. URL: https://doi.org/10.1126%5C%2Fscience.aab3782.

[10] Fiona M. Hollands and Devayani Tirthali. “MOOCs: Expectations and reality”. In: *Center for Benefit-Cost Studies of Education, Teachers College, Columbia University* 138 (2014).

[11] Rene F. Kizilcec, Mar Perez-Sanagustin, and Jorge J. Maldonado. “Self-regulated learning strategies predict learner behavior and goal attainment in Massive Open Online Courses”. In: *Computers & Education* 104 (Jan. 2017), pp. 18–33. DOI: 10.1016/j.compedu.2016.10.001. URL: https://doi.org/10.1016%5C%2Fj.compedu.2016.10.001.

[12] René F. Kizilcec, Andrew J. Saltarelli, Justin Reich, and Geoffrey L. Cohen. “Closing global achievement gaps in MOOCs”. In: *Science* 355.6322 (Jan. 2017), pp. 251–252. DOI: 10.1126/science.aag2063. URL: https://doi.org/10.1126%5C%2Fscience.aag2063.

[13] Kenneth R. Koedinger, Jihee Kim, Julianna Zhuxin Jia, Elizabeth A. McLaughlin, and Norman L. Bier. “Learning is Not a Spectator Sport”. In: *Proceedings of the Second (2015) ACM Conference on Learning @ Scale*. ACM, Mar. 2015. DOI: 10.1145/2724660.2724681. URL: https://doi.org/10.1145%2F2724660.2724681.

[14] Kenneth R. Koedinger, Elizabeth A. McLaughlin, Julianna Zhuxin Jia, and Norman L. Bier. “Is the doer effect a causal relationship?” In: *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge*. ACM, Apr. 2016. DOI: 10.1145/2883851.2883957. URL: https://doi.org/10.1145%2F2883851.2883957.

[15] Daphne Koller. *What we’re learning from online education*. 2012. DOI: 10.1037/e619852012-001. URL: https://doi.org/10.1037%2Fe619852012-001.

[16] Rita Kop. “The challenges to connectivist learning on open online networks: Learning experiences during a massive open online course”. In: *The International Review of Research in Open and Distributed Learning* 12.3 (Mar. 2011), p. 19. DOI: 10.19173/irrodl.v12i3.882. URL: https://doi.org/10.19173%2Firrodl.v12i3.882.

[17] Steven Leckart. “The Stanford education experiment could change higher learning forever”. In: *Wired* (2012). URL: https://www.wired.com/2012/03/ff_aiclass/.

[18] Allison Littlejohn and Colin Milligan. “Designing MOOCs for professional learners: Tools and patterns to encourage self-regulated learning”. In: *eLearning Papers* 42 (2015).

[19] Colin Milligan, Allison Littlejohn, and Anoush Margaryan. “Patterns of engagement in connectivist MOOCs”. In: *Journal of Online Learning and Teaching* 9.2 (2013), pp. 149–159.

[20] Sugata Mitra. “The Future of Learning”. In: *Proceedings of the Third (2016) ACM Conference on Learning @ Scale*. ACM, Apr. 2016. DOI: 10.1145/2876034.2876053. URL: https://doi.org/10.1145%2F2876034.2876053.

[21] Robert Murphy, Larry Gallagher, Andrew E. Krumm, Jessica Mislevy, and Amy Hafter. “Research on the use of Khan Academy in schools: Research brief”. In: *SRI International* (2014).

[22] Alyssa Napier, Elizabeth Huttner-Loan, and Justin Reich. “From online learning to offline action”. In: *Proceedings of the Fifth Annual ACM Conference on Learning at Scale*. ACM, June 2018. DOI: 10.1145/3231644.3231674. URL: https://doi.org/10.1145%2F3231644.3231674.

[23] Justin Reich. *EdTech start-ups and the Curse of the Familiar*. Nov. 2013. URL: http://www.edtechresearcher.com/2013/11/edtech_start-ups_and_the_curse_of_the_familar/.

[24] Justin Reich. *Failure to Disrupt*. Harvard University Press, Oct. 2020. DOI: 10.4159/9780674249684. URL: https://doi.org/10.4159%2F9780674249684.

[25] Justin Reich. *Failure to Disrupt Book Club with Cristina and Neil Heffernan*. 33. TeachLab with Justin Reich.

[26] Justin Reich. “Rebooting MOOC Research”. In: *Science* 347.6217 (Jan. 2015), pp. 34–35. DOI: 10.1126/science.1261627. URL: https://doi.org/10.1126%2Fscience.1261627.

[27] Justin Reich. “Will computers ever replace teachers.” In: *The New Yorker: Elements Blog* (2014). URL: https://www.newyorker.com/tech/annals-of-technology/will-computers-ever-replace-teachers.

[28] Justin Reich and Mizuko Ito. “From good intentions to real outcomes: Equity by design in learning technologies”. In: *Digital Media and Learning Research Hub* (2017).

[29] Mitchel Resnick, John Maloney, Andres Monroy-Hernandez, Natalie Rusk, Evelyn Eastmond, Karen Brennan, Amon Millner, Eric Rosenbaum, Jay Silver, Brian Silverman, and Yasmin Kafai. “Scratch”. In: *Communications of the ACM* 52.11 (Nov. 2009), pp. 60–67. DOI: 10.1145/1592761.1592779. URL: https://doi.org/10.1145%2F1592761.1592779.

[30] Jose A. Ruiperez-Valiente, Ruth Cobos, Pedro J. Munoz-Merino, Alvaro Andujar, and Carlos Delgado Kloos. “Early Prediction and Variable Importance of Certificate Accomplishment in a MOOC”. In: *Digital Education: Out to the World and Back to the Campus*. Springer International Publishing, 2017, pp. 263–272. DOI: 10.1007/978-3-319-59044-8_31. URL: https://doi.org/10.1007%2F978-3-319-59044-8_31.

[31] William A. Sandoval and Philip Bell. “Design-Based Research Methods for Studying Learning in Context: Introduction”. In: *Educational Psychologist* 39.4 (Dec. 2004), pp. 199–201. DOI: 10.1207/s15326985ep3904_1. URL: https://doi.org/10.1207%5C%2Fs15326985ep3904_1.

[32] George Siemens. “Connectivism: A learning theory for the digital age”. In: *Foundations of Educational Technology: Interdisciplinary Approaches to Educational Technology* (2016), pp. 2014–214.

[33] John Sweller. “Cognitive Load Theory”. In: *Psychology of Learning and Motivation*. Elsevier, 2011, pp. 37–76. DOI: 10.1016/b978-0-12-387691-1.00002-8. URL: https://doi.org/10.1016%5C%2Fb978-0-12-387691-1.00002-8.

[34] Edward L. Thorndike. “The Principles of Teaching Based on Psychology”. In: *The Elementary School Teacher* 6.8 (Apr. 1906), pp. 440–440. DOI: 10.1086/453574. URL: https://doi.org/10.1086%2F453574.

[35] Michael Trucano. *The Matthew Effect in educational technology.* 2013. URL: https://blogs.worldbank.org/edutech/matthew-effect-educational-technology.

[36] Yuan Wang, Ryan S. J. d. Baker, and Luc Paquette. “Behavioral predictors of MOOC post-course development”. In: *Proceedings of the Workshop on Integrated Learning Analytics of MOOC Post-Course Development*. 2017.

[37] Jacob Whitehill, Kiran Mohan, Daniel Seaton, Yigal Rosen, and Dustin Tingley. “MOOC Dropout Prediction”. In: *Proceedings of the Fourth (2017) ACM Conference on Learning @ Scale*. ACM, Apr. 2017. DOI: 10.1145/3051457.3053974. URL: https://doi.org/10.1145%2F3051457.3053974.

[38] Cheng Ye and Gautam Biswas. “Early Prediction of Student Dropout and Performance in MOOCs using Higher Granularity Temporal Information”. In: *Journal of Learning Analytics* 1.3 (Dec. 2014), pp. 169–172. DOI: 10.18608/jla.2014.13.14. URL: https://doi.org/10.18608%2Fjla.2014.13.14.