Computation, Constructivism, and Curriculum Design

Shayan Doroudi, Carnegie Mellon University, shayand@cs.cmu.edu

Abstract: Existing methods for automating curriculum design have had limited impact over the past fifty years. I propose two ways to potentially get around this limitation: developing adaptive content selection policies that are robust to different conceptions of student learning and taking an orthogonal approach to automated curriculum design that uses learner-generated solutions to help students learn.

Vision

From the mechanical teaching machines of the early twentieth century to intelligent tutoring systems and the wave of massive open online courses (MOOCs) in recent years, many have been motivated by the dream of personalized, adaptive instruction for all students. To achieve this dream, learning scientists and educational technology researchers have largely focused on rule-based systems that rely on extensive domain and psychology expertise. To do adaptive content selection, these systems use simple forms of rule-based AI (possibly combined with constrained machine learning algorithms). While this approach has led to the development of successful intelligent tutoring systems with high quality content, (1) such systems use a very limited form of adaptive content selection, and (2) developing such systems can be very costly. In contrast, some researchers are now starting to apply black box machine learning algorithms to do adaptive content selection. However, in my dissertation I will show through a comprehensive literature review that these approaches have had relatively limited impact.

Instead, I hope to demonstrate that combining insights from both approaches can help in automating curriculum design. In particular, I focus on three aspects of automated curriculum design: content creation, content curation, and adaptive content selection. I propose a number of methods for impactful, cost-effective automated curriculum design that combine machine learning, human computation, and principles from the learning sciences.

First, I will describe how reasoning about model mismatch (i.e., the fact that our statistical models of student learning do not accurately describe student learning) can help point out limitations in existing approaches [Doroudi and Brunskill, 2017] and help in creating more robust adaptive content selection policies [Doroudi et al., 2017].

Second, I will show experiments that demonstrate how we can leverage the work that students naturally do to create new content in a cost-effective way [Doroudi et al., 2016]. In doing so, I will take motivation from the constructivist philosophy of education, whereby I view learner-generated solutions as being a projection of students' constructions on the written plane, which can then be used to inform other students as they construct their own understandings.

Third, I propose to demonstrate how using machine learning (in particular, multi-armed bandits) can help curate the best content among a pool of learner-generated solutions that continues to grow over time.

Finally, I propose to show how we can use learning science principles to constrain the search for good content selection policies. In particular, I hope to show that constraining adaptive content selection algorithms with insights from the expertise reversal effect can help improve upon strictly black box approaches to adaptive content selection that disregard what we know about student learning.

References

Shayan Doroudi and Emma Brunskill. The misidentified identifiability problem of Bayesian Knowledge Tracing. In Educational Data Mining, pages 143–149. International Educational Data Mining Society, 2017.

Shayan Doroudi, Ece Kamar, Emma Brunskill, and Eric Horvitz. Toward a learning science for complex crowdsourcing tasks. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, pages 2623–2634. ACM, 2016.

Shayan Doroudi, Vincent Aleven, and Emma Brunskill. Robust Evaluation Matrix: Towards a more principled offline exploration of instructional policies. In Proceedings of the Fourth (2017) ACM Conference on Learning@ Scale, pages 3–12. ACM, 2017.