

SPEDS: A Taxonomy for Crowdsourcing in Education

Piotr Mitros, edX, pmitros@edx.org, Juho Kim, MIT, juhokim@mit.edu

Abstract: Over the past three years, we have run a series of experiments with having crowds of people create educational content. We present a taxonomy for crowdsourcing in education based on five elements: (1) Source (learner, instructor, crowd, etc.) (2) Pedagogical content knowledge required (3) Effort and complexity of task (4) Domain knowledge required (5) Amount of structure provided. We will show how this taxonomy helps guides the design of effective crowdsourcing strategies in educational settings.

Keywords: crowdsourcing, education, learnersourcing, MOOC, learning-at-scale

Introduction

There are 4,706 degree-granting institutions in the United States. An introductory course such as physics is taken by over one million students annually. This gives 3-4 orders of magnitude repetition and inefficiency on tasks such as creation of lectures and assessments, and 6 orders of magnitude on per-student tasks such as grading or tutoring. International numbers are approximately an order of magnitude greater. Open educational resources (OER) and at-scale learning organizations such as edX attempt to improve the quality of education by reducing those inefficiencies, providing higher-quality resources than previously possible, and freeing instructor time to allow for more time for student-instructor interaction.

Creating evidence-based, pedagogically effective resources is expensive. Some of the more effective technological resources for K-12 courses, with high-quality, evidence-based, interactive, adaptive, accessible, internationalized, and engaging content, cost millions of dollars to create. With tens of thousands of courses available at a university level, creating a repository of content at this level of quality would cost in the tens of billions of dollars. As a result, over the past three years, edX has run a series of experiments around sourcing such content from students, instructors, and crowds. Such techniques could potentially allow us to create better course resources at much lower cost, dramatically improving quality of tertiary education.

In the process of running these experiments, we came to several conclusions. First of all, there is a large body of individuals willing to help with course creation – students, instructors, and crowds. Such contributions can dramatically improve the quality of courses, including advanced and esoteric ones. Unfortunately, many crowdsourcing projects also failed primarily either due to lack of background of the participants, or lack of structure and guidance. In order to help understand success criteria, we have developed a taxonomy for describing such projects. The taxonomy has five axes:

- **Source:** This helps understand who is creating the content. For example, across experiments, this may have been current learners, alumni, crowds (as through Amazon Turk), and teachers. These may be from a MOOC or residential course. Furthermore, within those groups, participants were sometimes pre-filtered.
- **Pedagogical content knowledge:** Tasks require different levels of pedagogical content knowledge, as well as background in teaching and learning.
- **Effort and complexity of task:** Motivating learners to put in small amounts of time, as when contributing a hint is different from motivating them to contribute substantial time.
- **Domain knowledge required:** Many tasks, such as explaining relevance of information, or tagging information with concepts and learning objectives, requires a high level of domain knowledge.
- **Structure provided:** Crowdsourcing generally works better with a high level of guidance.

As we will show, as the scale of classes grows, as in MOOCs, these do not necessarily limit our ability to crowdsource content. There is a broad demographic of learners, including domain-experts and instructors (Ho, et. al. 2015). However, sourcing content complex on a large number of axes requires more nuanced approaches in order to be successful.

Successes: Learnersourcing Remediations

The simplest content to source from learners have been remediations for other learners. By the taxonomy:

1. **Source:** Current learners. Since there are many more learners, we only need low participation rates.

2. **PCK:** Learners fundamentally have a high level of PCK for creating remediations – they are still aware of what is difficult about those problems.
3. **Effort:** Small. Remediations are small, and relatively independent
4. **Domain knowledge required:** Learner-level.
5. **Structure provided:** In most cases, the task of providing remediations can be highly structured and guided. Learners can know exactly what to do.

We had a series of experiments in sourcing remediations. All were highly successful. Building on the success of AIQUS, a student-run Q&A forum in the Stanford AI Course, we used Askbot, an open source community question-and-answer system in MITx 6.002x, a MOOC on edX (Mitros, et. al. 2013). Students who asked questions would receive an answer 92% of the time, with a median response time of 12 minutes. The quality of responses were above what instructors in a typical class would provide. The Q&A forum was also a repository of answers to common questions. If a question has been answered, that answer is usually found rather than having a question re-asked – students read 290 times as many threads as they created. As a result, there is a very large number of potential contributors for any given question. For more immediate remediation, in a course in structures, we allowed students who answered a question incorrectly and later correctly to contribute a hint for future learners who made the same error. This allowed us to create a body of hints similar to what might be found in a commercial system, such as MasteringPhysics, but substantially outclassing such a system for number of remediations (Mitros & Sun, 2014). In order to give remediations with less scaffolding, we have experimented with allowing students to contribute links to relevant resources in a course in computer science, which we then recommend to other students. This, too, was highly successful (Li & Mitros, 2015).

Difficulty: High Expert-Novice Gap Content

We had a series of experiments where we tried to source content with a greater expert-novice gap from learners. In one experiment, we took a group of residential MIT 6.002 alumni and asked them to tag problems from 6.002x with learning objectives. Students would tag on surface features, rather than on core concepts. This replicated a well-known result in educational cognitive science – beginners are more likely to focus on surface features (“this problem has a diode”) while experts, on core concepts (“This problem introduces non-linear devices”) (Chi, et. al. 1981).

In another experiment, we provided students with a wiki, seeded with lecture outlines. We hoped students would be able to collaboratively create an open access textbook for 6.002x. Students were able to contribute – substantially – but the result was a shared set of course notes. Again, expert-novice literature states that novices can explain knowledge well, but are not capable of contextualizing it sufficiently to explain it in ways with lead to high levels of transfer (Hinds & Patterson & Pfeffer. 2001). In addition, students were not provided with a high level of guidance about how to contribute (most contributors were unaware of our ultimate goal). Unstructured crowdsourcing tends to be less successful than well-structured crowdsourcing.

Overcoming Limitations to Crowdsourcing Complex Content

Creating many portions of a course requires a high level of domain expertise, PCK, technical background, and effort. Our effort in remediations demonstrated that we could substantially shift the quality of courses with crowdsourcing, but substantially shifting economics requires creation of complex content – videos, assessments, course texts, and similar. We tried two experiments to generate this type of complex content.

In the first experiment (Cormier, et. al. 2014), we worked with a group of experts. We organized a course in how to teach physics, and brought together a pool of physics educators, education researchers, technology experts, and similar. They were asked to create digital course content while studying ed-tech and physics education research. This was a small pilot (4 weeks, limited participant pool) in preparation for a larger roll-out. This appeared to be a viable approach for a basic course, where there are thousands of instructors available. However,

technical limitations with our authoring tools prevented us from moving beyond the pilot to fully validate the concept.

For more esoteric university courses, there is not a large pool of potential instructors, and we must rely on learners. We ran an experiment where we asked learners in 6.002x to create assessments for the course. This was largely successful, but required a complex approach involving: (1) pre-selecting top alumni (2) training those alumni in engineering pedagogy (3) community review and feedback on resources (4) instructor review of resources. This was a complex endeavor, but it could be fully automated if done across hundreds of courses (Mitros, 2015). Note that this approach is specific to a MOOC context – MOOCs have very heterogeneous student bodies, including instructors trying to learn new pedagogical techniques, experts review content, etc. (Ho, et. al. 2015). It is likely this approach would not have worked in a traditional classroom setting without those experts.

Conclusions and implications

Through a series of studies in educational crowdsourcing conducted primarily in MOOCs, we found that course quality and economics could be substantially improved through the involvement of crowds in course creation. However, naive approaches will yield mixed results. Course content creation requires a complex and diverse skillset, and when there is a mismatch between the expertise of participants and the expertise required, the contributed content is likely to be of sub-par quality.

However, in such cases, less naive approaches can yield good results. As a procedure, identifying the principal gaps is key to being able to create an effective crowdsourcing strategy (our taxonomy provides good guidance here). From there, there are many approaches one can take to resolving those gaps. A typical MOOC might have on the order of 10,000 active participants, including thousands of subject-matter experts, instructors, and similar. While only a minority of these will be willing to contribute, this still leaves a sufficient pool for most crowdsourcing purposes. It is also possible to fill the gaps through specific training. In most cases, participants who are likely to contribute also appreciate the additional instruction. It is also possible to create structure around the task which provide guidance to how best to contribute, permit the best contributions to bubble to the top, and gives the community means to improve contributions. Once expertise gaps are identified, it is generally straightforward to identify which strategies are likely to be effective.

References

- Chi, M. T. H., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5: 121-152.
- Cormier* & Mitros* & Pritchard*. (2014). Teaching Physics at Scale. *Physics Education Research Conference*.
- Hinds, P. J., Patterson, M., & Pfeffer, J. (2001). Bothered by abstraction: the effect of expertise on knowledge transfer and subsequent novice performance. *Journal of applied psychology*, 86(6), 1232.
- Ho & Chuang & Reich & Coleman & Whitehill & Northcutt & Williams & Hansen & Lopez & Peterson. (2015). HarvardX and MITx: Two Years of Open Online Courses Fall 2012-Summer 2014. *Social Sciences Research Network*.
- Li, Mitros (2015). Learnersourced Recommendations for Remediation. *IEEE International Conference on Advanced Learning Technologies* (accepted).
- Mitros (2015). Learnersourcing of Complex Assessments. *ACM Learning at Scale*.
- Mitros & Afridi & Sussman & Terman & White & Fischer & Agarwal (2013). Teaching electronic circuits online: Lessons from MITx's 6.002x on edX. *IEEE International Symposium on Circuits and Systems*.
- Mitros & Sun. (2014). Creating Educational Resources at Scale. *IEEE International Conference on Advanced Learning Technologies*.

Acknowledgments

We thank the students and teachers who helped contribute to our courses. Without them, this line of work would be impossible.