How Can Crowdsourcing Help Individuals Learn?

Kurt Luther
HCI Institute
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213 USA
kluther@cs.cmu.edu

Steven Dow
HCI Institute
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213 USA
spdow@cs.cmu.edu

Aniket Kittur
HCI Institute
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213 USA
nkittur@cs.cmu.edu

Submitted to CSCW '14 Workshop on
Designing Futures for Peer-to-Peer Learning

Abstract
In this paper, we consider how crowdsourcing can augment an individual's learning process. We first describe the problem space, identifying the promise and limitations of crowdsourcing for supporting learning. We then propose a vision of crowdsourced support for learning across a range of situations that considers the needs of both the learner and the crowd. Finally, we report on two of our projects in this space, CrowdCrit and CrowdLines, that take steps towards realizing this vision of crowd-augmented learning.

Introduction and problem space
Crowdsourcing has proven to be a powerful form of social computing, enabling new types of problem solving and creative production across a range of domains, from citizen science to product development to entertainment and the arts. In this paper, we argue that crowdsourcing can also provide powerful new learning experiences for individuals.

Previous work has demonstrated the benefits of social learning, and ways that online communities can enhance these educational benefits [2,5]. Crowdsourcing uses new incentives and structures that can provide many of the same benefits with increased control and flexibility. For example, crowds can scale dynamically, supporting large numbers of learners.

Technology allows geographically distributed crowds to connect with learners, enabling learning experiences beyond traditional co-located settings like classrooms. Crowds can offer high availability so that interactions with learners can occur rapidly or even in real time. Crowds can also represent diverse backgrounds and skill sets that enrich learners' own abilities.

However, using crowdsourcing to promote learning also raises a number of logistical and ethical concerns that must be considered. On paid crowdsourcing platforms like Amazon Mechanical Turk, crowd workers are usually employed for unskilled microtasks, such as labeling images. Attempts to use the crowd to teach complex skills or domain-specific knowledge must address the fact that crowds are mostly comprised of novices. Even if a subset of workers possesses the required knowledge, this limited talent pool will not scale well. Finally, most crowd platforms are oriented around financial incentives and favor employers over workers [4]. The systems we build to support learning should provide workers with fair compensation, meaningful work, and protection from abuse.

Our vision

We envision a future where *crowds can help individuals learn*. In this future, the experience for both crowds and learners is rewarding. A range of learning situations would be supported, from students working on class assignments to scientists exploring an unfamiliar research area to individuals learning a new hobby. The experience would be personalized to the learner's background and goals, and would dynamically adjust as the learner gains new knowledge, creating a "flow" state that balances challenge and accomplishment [3]. Crowds would provide both technical and social support, depending on the learner's needs. Learners could access the crowd from anywhere around the world, at any time of day. Learners could either participate in a reciprocal market (earning credits by joining the crowd) or pay a reasonable fee to receive crowd assistance.

The crowd's experience in supporting learners would also be meaningful and rewarding. Crowd workers would have the opportunity to learn as well, building up their skill sets and completing more complex, creative tasks. If they choose to be financially compensated, the payment would be prompt and fair. Workers would be able to see how their efforts assist the learner and how their efforts aggregate with other workers.

Making the vision a reality

We are currently involved in two projects that take small but important steps towards achieving the vision described above. The first, CrowdCrit, investigates how crowds can provide individuals with high-quality feedback on creative work, focusing on the domain of visual design. The second project, CrowdLines, investigates how crowds can help an individual understand an unfamiliar body of knowledge.

CrowdCrit

This project, a collaboration with Björn Hartmann, Maneesh Agrawala, Wei Wu, and Amy Pavel at the University of California, Berkeley and Jay Tolentino at the University of California, Irvine, seeks to use crowdsourcing to provide high-quality feedback to visual designers. Many people without formal design training regularly engage in "everyday design," such as creating posters, flyers, and slide decks. Their designs may not be effective, and they lack access to sources of

Layout
Good alignment
Ueven margins
Poor cropping
Readability
No spelling errors
Poor kerning
Background lacks contrast
Simplicity
Simple and clean
Overuse of images
Lacks white space
Emphasis
Strong focal point
Lacks hierarchy
False proximity
Balance
Good use of symmetry
Lacks balance
Lacks movement
Consistency
Good repetition
Element disrupts unity
Poor consistency
Appropriateness
Reaches intended audience
Mixed messages
Inappropriate message

Table 1. Design principles and sample critique statements (positive and negative).

high-quality feedback to help them learn and improve. Crowdsourcing offers an appealing potential solution because crowds are fast and scalable, but most crowd workers also lack design expertise. Thus, this project has two interleaved learning goals. First, we want to help crowd workers learn the process and language of design critique so they can provide high-quality feedback. Second, we want to help designers improve by giving them access to high quality crowd critique. Our solution takes the form of a web-based critique system called CrowdCrit.

To help the crowd develop critique skills, we employ a learning technique called *scaffolding*, which structures the learning process to help people accomplish more than they could unaided [2]. By consulting design textbooks, we developed a set of 70 critique statements across 7 high-level design principles (Table 1). Using a critique interface, crowd workers from Amazon Mechanical Turk can view a design and describe a strength or weakness by selecting an appropriate critique statement from the list. The worker can elaborate by annotating an area of the design and providing details as an open-ended comment. These features scaffold the characteristics of a high-quality critique: specific, conceptual, and actionable.

After multiple crowd workers critique the design, the designer reviews these critiques using an aggregation interface. An interactive visualization provides an overview of all the critiques, organized by design principle and color-coded as strengths or weaknesses. The designer can drill down on a principle to see the individual critiques, text comments, and graphical annotations for each critique statement. This interface allows the designer quickly identify the strengths and weaknesses in the design, but also reflect on their relationship to broader design principles and suggestions for improvement.

To evaluate CrowdCrit, we organized a poster design contest in which 14 designers received crowd critiques between iterations. We found that designers acted on many of the critiques to modify their designs, and designs with more negative critiques were more likely to show improvement. These results suggest that scaffolding can help crowd workers learn complex information and perform more creative tasks. They also suggest that aggregated crowd feedback can help novice individuals reflect on and improve creative efforts. In future work, we are experimenting with ways to improve and measure what crowds and individuals learn while using CrowdCrit.

CrowdLines

CrowdLines explores how crowds could help an individual make sense of an unfamiliar body of knowledge. This project focuses on the context of scholars seeking to understand a new research area. With academic publications growing exponentially, it is increasingly difficult for scholars to keep up to date in their own areas of specialization. It is even more challenging for them to gain familiarity with areas outside of their expertise, leading to redundant projects, oversights of related work, and missed opportunities for collaboration.

We propose that crowdsourcing can help scholars more effectively explore these unfamiliar research areas, make sense of them, and identify interesting connections. Our approach is to use crowds to generate dynamic, personalized, richly annotated outlines of the

scholar's topic of interest, called CrowdLines. Crowds work in parallel with the learner (scholar), leveraging the massive amounts of information available online to find, prioritize, summarize, and aggregate existing high-quality materials into an interactive overview of the topic. As the learner dives deeper into particular subtopics of interest, the crowd responds by elaborating those sections of the CrowdLine with more specific, detailed information and references. CrowdLines can be reused in future sensemaking endeavors by the original learner or others interested in similar topics.

The CrowdLines project is a work-in-progress. Our current research focuses on two major threads. First, what are the most effective strategies for crowd workers to synthesize multiple, potentially conflicting perspectives on the most important issues within a research topic? We are presently experimenting with two approaches, mapping and clustering, based on recent work in crowdsourcing and databases research (e.g. [1]). Second, what are the most effective ways for the learner to interact with the crowd? Our goal is to actively engage the learner while using the crowd to extend and deepen the learner's understanding of the topic. Currently, we are exploring theories of leadership and improvisation as potential frames for the interaction patterns between the learner and the crowd. Our future plans involve developing and evaluating a prototype of the CrowdLines system. We expect this system and evaluation will enrich our understanding of ways that crowdsourcing can help people learn and inspire new forms of creative social computing.

Biography

Kurt Luther is a postdoctoral fellow in the HCI Institute at Carnegie Mellon University. His research investigates how social computing can support creativity, focusing on key issues such as coordination, expertise, and ethics. At the workshop, he hopes to contribute theoretical and practical knowledge surrounding how to leverage crowdsourcing approaches and technologies to support learning and creativity. He hopes to share his research experiences in this domain and learn how others are using social computing in educational contexts. He is especially interested in unpacking the ethical issues involved in crowdsourced learning and gaining a deeper appreciation of the educational theories and evaluation strategies used to conduct this type of research.

References

1. André, P., Kittur, A., and Dow, S.P. Crowd Synthesis: Extracting Categories and Clusters from Complex Data. *Proceedings of CSCW 2014* (2014).
2. Bransford, J.D., Brown, A.L., and Cocking, R.R., eds. *How people learn: brain, mind, experience, and school*. National Academy Press, Washington, D.C., 2000.
3. Csikszentmihalyi, M. *Flow: The Psychology of Optimal Experience*. Harper Perennial, 1991.
4. Irani, L.C. and Silberman, M.S. Turkopticon: Interrupting Worker Invisibility in Amazon Mechanical Turk. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2013), 611–620.
5. Resnick, M., Maloney, J., Monroy-Hernández, A., et al. Scratch: Programming for All. *Commun. ACM 52*, 11 (2009), 60–67.