MIT Open Access Articles

Crowdsourcing and Human Computation: Systems, Studies and Platforms

Citation	Michael Bernstein, Ed H. Chi, Lydia Chilton, Bjorn Hartmann, Aniket Kittur, and Robert C. Miller. 2011. Crowdsourcing and human computation: systems, studies and platforms. In PART 2 ----------- Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems (CHI EA '11). ACM, New York, NY, USA, 53-56.
As Published	http://dx.doi.org/10.1145/1979742.1979593
Publisher	Association for Computing Machinery (ACM)
Version	Author's final manuscript
Citable link	http://hdl.handle.net/1721.1/79668
Terms of Use	Creative Commons Attribution-Noncommercial-Share Alike 3.0
Detailed Terms	http://creativecommons.org/licenses/by-nc-sa/3.0/

Crowdsourcing and Human Computation:
Systems, Studies and Platforms

Michael Bernstein
MIT CSAIL
32 Vassar St.
Cambridge, MA 02139
msbernst@mit.edu

Ed H. Chi
Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA
ed.chi@parc.com

Lydia Chilton
University of Washington
AC101 Paul G. Allen Center
185 Stevens Way
Seattle, WA 98195
hmslydia@cs.washington.edu

Björn Hartmann
University of California, Berkeley
629 Soda Hall
Berkeley, CA 94720
bjoern@eecs.berkeley.edu

Aniket Kittur
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
nkittur@cs.cmu.edu

Robert C. Miller
MIT CSAIL
32 Vassar St.
Cambridge, MA 02139
rcm@mit.edu

Abstract

Crowdsourcing and human computation are transforming human-computer interaction, and CHI has led the way. The seminal publication in human computation was initially published in CHI in 2004 [1], and the first paper investigating Mechanical Turk as a user study platform has amassed over one hundred citations in two years [4]. However, we are just beginning to stake out a coherent research agenda for the field. This workshop will bring together researchers in the young field of crowdsourcing and human computation and produce three documents: a research agenda for the field, a vision for ideal crowdsourcing platforms, and a group-edited bibliography. These documents will be publically available on the web and maintained by the workshop organizers.

Keywords
Crowdsourcing, human computation, workshop, bibliography, platform.

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

Introduction

Crowdsourcing and human computation are transforming human-computer interaction, and CHI has led the way. The seminal publication in human computation [1] was initially published in CHI in 2004, and the first paper investigating Mechanical Turk as a user study platform has amassed over one hundred citations in two years [4]. Crowdsourcing and human computation research, so far, has involved paying or motivating many people to contribute work to a research goal – which might be a user study, a novel interface, or a complex cognitive task. We are just beginning to learn what's possible, however, and we need to stake out a research agenda for the field.

Outside of CHI, last year saw over a dozen computer science conferences with crowdsourcing workshops, including SIGIR, CVPR, NACL, NIPS, and Ubicomp. Most of these workshops target crowdsourcing primarily as a tool -- e.g., labeling datasets for learning -- rather than advancing the state of the art in understanding and supporting large-scale interaction among people. Since human computation started in the CHI community, the next wave of innovations should be expected to come from CHI as well.

This workshop will bring together members of the CHI community who have demonstrated interest in crowdsourcing and human computation, in order to lay out a framework for continued innovation. We are interested in researchers who are:

1. Studying how to use crowdsourcing as a tool in user-centered design and HCI research, such as user testing, formative evaluation, or controlled experiments;

2. Building and studying systems powered by human computation and the contributions of crowds;

3. Studying crowdsourcing platforms like Amazon Mechanical Turk, and interested in how to engineer better platforms.

In addition to community building, we want to produce three concrete deliverables:
- a research agenda for the field
- a set of requirements for an ideal crowdsourcing platform for HCI work
- a group-edited bibliography.

These deliverables will be posted on a web site and maintained afterwards by the workshop organizers.

Workshop Plan and Goals

This one-day workshop will also involve additional activities organized remotely both before and after the workshop date.

PREPARATORY ACTIVITIES

Position statements.
Potential participants will submit a 2-page abstract about their experience in and interests in crowdsourcing and human computation. Participants will be selected because of their unique experience of working in crowdsourcing and human computation and different visions of how the field should progress.

Bibliography.
Crowdsourcing and human computation are growing very quickly and have taken root in many different conferences, making it difficult to get a full picture of the field, particularly for new researchers. The conference organizers will set up an open, web-accessible and searchable bibliography of human computation work (e.g., on Mendeley.com). Accepted participants will be asked to contribute 3 - 5 favorite papers to this online archive in the weeks before the workshop. Similar archives have become extremely valuable resources for academics in other subfields (e.g., [2, 3, 5]).

WORKSHOP ACTIVITIES

Introductory talks.
In the first morning session on the workshop day, participants will give short (2-minute) talks introducing themselves and their interests. Talks will be Ignite-style: slides automatically advancing every 15 seconds to keep things lively. Participants might discuss work they've done in crowdsourcing, problems they've encountered in existing human computation systems, ideas for future directions, and aspects of their position statements.

State of the field.
The second morning session will be a breakout session around existing research. Researchers will form groups by broad topic area and draw out interesting themes and challenges in each area.

For example, one group might consider design dimensions for the space of crowd-powered systems:
- number of contributors,
- number of requesters,
- amount of latency that can be tolerated,
- ways to motivate contributors,
- expertise demanded from contributors.

Another group might consider evaluation:
- What properties of a crowdsourcing system or platform need to be measured?
- What is correct reporting for a crowdsourced study?
- How big should the sample be?
- What are best practices, and what techniques will lead to bias?

Another group might consider boundaries:
- What crowdsourcing research is in the purview of CHI?
- What do overlapping subfields investigating Wikipedia, Delicious, Facebook and Twitter have to contribute to crowdsourcing research?

Platform requirements.
The first afternoon session will be a breakout session to brainstorm the ideal crowdsourcing platform. Amazon's Mechanical Turk is a popular platform for this research, but it is far from perfect. What would the ultimate Mechanical Turk look like, if we could design it ourselves? What about a platform for prototyping systems that use an *unpaid* crowd, like the next Wikipedia or Ushahidi? Questions considered in this session might include:
- How can we design strong reputation systems for workers and requesters, while respecting privacy and preserving anonymity?
- How can we obtain demographic information, again respecting individual privacy?
- How can we select participant pools based on demographics?
- How can we run between-subjects studies with confidence?
- How can researchers join forces to prototype and advertise new crowdsourcing systems with motivations other than pay?

The outcome of this discussion would be a set of requirements that could be taken to Amazon or one of its competitors (such as CrowdFlower), or inspire third-party tools on top of existing platforms, or even motivate our own community to gather the resources required to build a new platform ourselves.

Research agenda.
The last session of the day will identify near-term research questions and longer-term grand challenges for the field. Where are we going, and how do we get there? What are important theoretical questions to ask? What problems would transform our field if solved? Conversely, which issues seem to be transient properties of existing systems or platforms and are not likely to stand the test of time?

Organization of Breakout sessions

The three breakout sessions (state of the field, platform requirements, and research agenda) will be organized as follows:

1. Seed discussions with broad topics. For example, for the research agenda discussion, we propose topics like "crowd-powered systems", "interfaces for crowd workers", "online experiments", and "ethical issues".

2. Participants choose a topic they are interested in discussing and cluster into small groups of 4-5 people.

3. Open-ended small group discussion on the topic, with groups recording their ideas on a collaborative online notepad such. Groups have 15-20 minutes.

4. Small groups present their ideas for general comment, and group comments are added to the collaborative notepad.

5. Later, the organizers synthesize the documents together.

FOLLOW-UP

Continuing community.
In addition to the research agenda and platform requirements, which will be published on the workshop's web site, the workshop will have continuing deliverables. The bibliography will be maintained and expanded as the field grows. Depending on participant interest, we will also explore creating a community blog to share summaries of publications, opinion pieces, experiments, and announcements of conferences and workshops.

References

[1]   von Ahn, L. and Dabbish, L. Labeling images with a computer game. I*CHI 2004*.

[2]   boyd, d. Bibliography of Research on Social Network Sites.
http://www.danah.org/researchBibs/sns.php

[3]   boyd, d. Bibliography of Research on Twitter & Microblogging.
http://www.danah.org/researchBibs/twitter.php

[4]   Kittur, A., Chi, E., Suh, B. Crowdsourcing User Studies With Mechanical Turk. *CHI 2008.*

[5]   Whitten, A. HCISec Bibliography
http://gaudior.net/alma/biblio.html