
Communicating Context With Microtasks

Niloufar Salehi

Stanford University
California, USA
niloufar@cs.stanford.edu

Ece Kamar

Microsoft Research
Redmond, WA USA
eckamar@microsoft.com

Jaime Teevan

Microsoft Research
Redmond, WA USA
teevan@microsoft.com

Shamsi T. Iqbal

Microsoft Research
Redmond, WA USA
shamsi@microsoft.com

Bio

Niloufar Salehi is a third year PhD student in computer science at Stanford University and a member of the Human-Computer Interaction group. Her research focuses on crowdsourcing, social computing, and collective action online. In particular, she is interested in designing and building systems that guide the efforts of a group of people towards a shared goal. She has previously collaborated with researchers and Turkers to build Dynamo, a system that enables Turkers to gather around shared issues and act on them together. She is currently working on designing mechanisms for efficient communication between crowd workers and requesters.

Designing Efficient Communication Mechanisms for Crowd Work

Crowdsourcing is primarily used to perform short, independent microtasks that anyone could do. However, there is a recent shift towards enabling more complex roles that handle interrelated parts and demand awareness of context. For instance, Cobi enables workers to view the current state of global constraints while performing tasks [2, 4], and Turkomatic [3] empowers workers to take responsibility in dividing tasks into subtasks, a process that had previously been done by the requestor or through automatic workflows (e.g. find-fix-verify [1]). As the crowd's contributions become more complex and

Paste the appropriate copyright/license statement here. ACM now supports three different publication options:

- ACM copyright: ACM holds the copyright on the work. This is the historical approach.
- License: The author(s) retain copyright, but ACM receives an exclusive publication license.
- Open Access: The author(s) wish to pay for the work to be open access. The additional fee must be paid to ACM.

This text field is large enough to hold the appropriate release statement assuming it is single-spaced in Verdana 7 point font. Please do not change the size of this text box.

Each submission will be assigned a unique DOI string to be included here.

interrelated, tasks can no longer be done independent of the worker, other tasks, or the information space surrounding the task; we refer to all of these external factors that influence a task's outcome as its *context*. For example, a crowd worker asked to write a paragraph about a research project might find it useful context to know the intended audience and have relevant terms defined.

It is not always possible to make the necessary context available before the tasks starts. For example, the requester described above who was looking for the crowd to write a paragraph about their research project may not know until they see the first draft that they would like it to use formal language. This type of context transfer is only possible through communication. Despite how important communication is for many tasks, context transfer is not well supported within current crowdsourcing systems such as Amazon Mechanical Turk. These platforms are limited to objective, short tasks with no context. In contrast, on expert crowdsourcing platforms like Upwork, requesters and workers engage in one-on-one, open-ended discussions using chat software and Skype. Requesting and integrating context is a challenging yet crucial part of an expert crowd worker's job. Back-and-forth communication is necessary to achieve mutual understanding in complex domains, but these open-ended discussions tend to be costly and time consuming, with important context often coming later in the process than when it is needed.

My research explores how to use structured microtasks to support communication related to a larger task, comparing different mechanisms based on the cost to the requester for providing context and the value of the

provided context to workers. We find different mechanisms are effective at different stages of a task, and that it is possible to transfer some of the effort of providing feedback from the requester to the crowd via question asking. These findings can be used to enable richer, more complex crowd work than currently seems possible. For example, we develop a system that creates written content from a list of bullet points using appropriately timed structured mechanisms for communicating with the crowd.

References

1. Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. "Soylent: a word processor with a crowd inside." In Proceedings of the 23rd annual ACM symposium on User interface software and technology, pp. 313-322. ACM, 2010.
2. Juho Kim, Haoqi Zhang, Paul André, Lydia B. Chilton, Wendy Mackay, Michel Beaudouin-Lafon, Robert C. Miller, and Steven P. Dow. "Cobi: A community-informed conference scheduling tool." In Proceedings of the 26th annual ACM symposium on User interface software and technology, pp. 173-182. ACM, 2013.
3. Anand Kulkarni, Matthew Can, and Björn Hartmann. "Collaboratively crowdsourcing workflows with turkomatic." In Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work, pp. 1003-1012. ACM, 2012.
4. Haoqi Zhang, Edith Law, Rob Miller, Krzysztof Gajos, David Parkes, and Eric Horvitz. "Human computation tasks with global constraints." In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 217-226. ACM, 2012.