

MobileWorks: A Mobile Crowdsourcing Platform for Workers at the Bottom of the Pyramid

Prayag Narula¹, Philipp Gutheim¹, David Rolnitzky¹, Anand Kulkarni², Bjoern Hartmann³

¹School of Information, ²Industrial Engineering and Operational Research, ³Computer Science Division
University of California, Berkeley

{prayag.gutheim,dr}@ischool.berkeley.edu, anandk@berkeley.edu, bjoern@cs.berkeley.edu

Abstract

Existing crowdsourcing markets are often inaccessible to workers living at the bottom of the economic pyramid. We present MobileWorks, a mobile phone-based crowdsourcing platform intended to provide employment to developing world users. MobileWorks provides human optical character recognition (OCR) tasks that can be completed by workers on low-end mobile phones through a web browser. To address the limited screen resolution available on low-end phones, MobileWorks divides documents into many small pieces and sends each piece to a different worker. An initial pilot study with 10 users over a two month period revealed that it is feasible to do basic OCR tasks using a simple mobile web-based application. We find that workers using MobileWorks average 120 tasks per hour at an accuracy rate of 99% using a multiple entry solution. In addition, users had a positive experience with MobileWorks: all study participants would recommend MobileWorks to friends and family.

Introduction

Amazon’s Mechanical Turk (AMT) has about 200,000 workers who reside primarily in the US and India. About 56% of these Mechanical Turk workers (“Turkers”) are from the United States and 36% from India. Among the Indian population, Turkers are often more educated, earn higher wages, and have a higher standard of living than the average Indian (Ipeiritis, 2010). Yet many Indians have limited English literacy and lack access to a desktop computer--the computer penetration in India in 2004 was only about 0.9% (Dwivedy and Mittal 2010)--making most crowdsourcing work inaccessible to a large proportion of the Indian population. Investigations into whether AMT could be used as an employment and development tool for workers at the bottom of the pyramid found that even workers with basic computer training could not effectively use the platform due to inaccessible task interfaces and cultural barriers (Khanna, 2010).

Copyright © 2011, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: A worker in India completes OCR tasks on his cell phone using MobileWorks.

In contrast to desktop computers, mobile phone penetration in India is very high (about 50%). Many of these phones are simple, “candy-bar” style phones capable of surfing the web. Moreover, the cost of mobile Internet for these basic cellphones has plummeted in India in recent years. In a formative survey, we found access to the mobile Internet is available at Rs.50 (\$1.10 US) a month in India. In many locations, three-day plans are available from Rs 15 (\$0.33 US). Because of the ubiquity of inexpensive cell phones, even people who earn less than \$2.00 US a day have access to the Internet. As a result, the mobile Internet is an cost-effective way to send microtasks to people at the bottom of the economic pyramid (Lehdonvirta, 2011).

In response to this opportunity, we are developing MobileWorks, a mobile microtask platform for developing countries (Figure 1). The goal of MobileWorks is to broaden participation in microtask markets by marginalized workers. A key research question is whether mobile interfaces can be made efficient enough so as to provide livable wages to workers, and at the same time provide organizations with data digitization needs a low-cost, accurate, and efficient solution. We investigate these questions through a pilot deployment of OCR tasks.

Related Work

Using cellphones to distribute microtasks has been explored before. TxtEagle, deployed in Kenya, used SMS text messages to provide tasks like audio transcription, local language translation and market research (Eagle, 2009). However, because they rely on text SMS as the medium of communication and are restricted in length, there are a limited number of tasks that can be accomplished through the use of simple text messages.

SamaSource (Thies, Ratan, & Davis, 2010) also seeks to empower marginalized workers. However, in contrast to TxtEagle, SamaSource does not use mobile crowdsourcing; instead they establish outsourcing centers in developing regions where workers are actively managed.

Human OCR has also been used on desktop-based crowdsourcing solutions to produce accurate results for blurry text (Little, Chilton, Goldman, & Miller, 2009). The authors compared two task flows: a parallel work flow in which different micro workers do the same OCR tasks independently, and an iterative work flow where each subsequent worker is shown answers from the previous workers. The authors reported that using iterative task flow it is possible to achieve 71% accuracy for even very blurry text.

Human OCR using MobileWorks

MobileWorks consists of three main components. First, a preprocessing stage takes scanned paper documents and generates many small OCR tasks for each document. Second, a user interface enables workers to perform OCR tasks. Third, a post-processing stage reassembles completed OCR tasks and ensures quality.

MobileWorks process

Because of the limited screen size on mobile phones, documents have to be chopped into small pieces of one or two words. Different workers digitize the pieces using the Mobileworks web application and submit it to the server. These smaller pieces are then put together to create a digitized copy of the document.

Quality

Quality is maintained using multiple entry. Each task is distributed to two workers until two of the answers match. If a worker provides an incorrect answer, her quality rating decreases. Conversely, a worker that provides a correct answer will see an increase in quality score.

Accuracy

We use the historical accuracy of the worker to model the future payment of the tasks she is assigned. Hence, for each worker, an a priori payment is a function of the task and her quality of the work. Since the cost of each task is calculated before the task is complete, the worker can view how much money she has made in real-time.

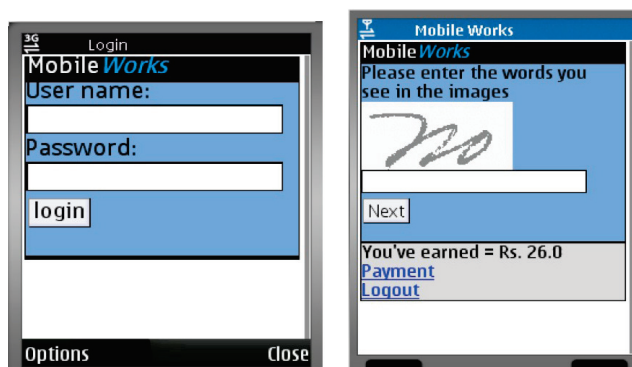


Figure 2: The MobileWorks web application user interface

Pilot Study

We tested a set of handwritten documents and scans from the stock page of 19th century newspapers. 10 workers were recruited in two locations: Dharavi, a slum in Mumbai, and in Rewari, a small town near New Delhi. The pilot was carried out over a period of two months, with a set of 1,500 images digitized by the group of workers. Each image was digitized twice, comprising a set of 30,000 tasks total.

The interface was minimal, as illustrated in Figure 2. In the first image (left), the user logs in using her username and password. The user can start doing tasks immediately after login.

Accuracy

We found that the overall accuracy of the workers, without considering multiple entry error detection, was about 89%. The accuracy was highest for the handwritten documents and lowest for 19th century newspaper stock data. This number should be taken as a lower bound on the accuracy possible with the system. Though we did not explicitly test the accuracy using multiple entry in the experiment, in cases where errors were independently distributed between workers, dual entry accuracy would be 98.79%, and for triple entry 99.89%. This suggests that by using multiple entry it is possible to provide very high quality crowdsourced human OCR work on simple mobile phones.

Efficiency

In our survey we found that participants were earning 20 - 25 Indian Rupees an hour (USD 0.55) while doing their regular job. While using MobileWorks, participants were able to complete 120 tasks per hour. Therefore, on an hourly basis the workers should be paid approximately 0.18 to 0.20 Indian Rupees per task to match their regular wages. We predict that the efficiency of the workers would increase as they become more competent at doing tasks on their mobile phone. In addition, improvements in wireless infrastructure will improve data transfer rates, decreasing the time it takes to complete each task.

Survey

We conducted a survey at the end of the experiment to evaluate the usability of the system. Overall, the response of the workers was positive. When asked to rate the usability of the system on a five point Likert Scale, ten out of ten users rated the usability above a four. Moreover, all of the users said that they were more than likely to recommend the system to their friends and family (four or above on a five point scale). Pilot participants said that the biggest advantage of the system was that the work could be done anywhere, at anytime of the day: “I could do the work and earn money while traveling to my regular job or even while watching television” said one of the participants.

Conclusion and Future Work

MobileWorks is a mobile web-based crowdsourcing platform which we used to conduct human OCR tasks. Future work will explore the possibility of doing other types of work using MobileWorks, such as audio transcription, same language subtitling and local language translation. Future work will also explore the possibility of providing near real-time response by allowing posting of high priority tasks to the front of the task queue.

References

- Bernstein, M. S., Little, G., Miller, R. C., Hartmann, B., Ackerman, M. S., Karger, D. R., et al. (2010). Soylent : A Word Processor with a Crowd Inside. *Artificial Intelligence*, 313-322.
- Dwivedy, M. and Mittal, R.K. (2010) Future trends in computer waste generation in India. *Waste management*, vol. 30, 2265-77.
- Eagle, N. (2009). txteagle : Mobile Crowdsourcing, *Internationalization, Design and Global Development*, 447-456.
- Ipeirotis, P. G. (2010). Analyzing the Mechanical Turk Market Place. *XRDS*, 17(2), 16-21.
- Khanna, S., Ratan A, Davis, J., Thies, W. (2010). *Evaluating and Improving the Usability of Mechanical Turk for Low-Income workers in India*, *ACM SigDev*
- Lehdonvirta, V., Enrkvist M. (2011). Knowledge map of the virtual economy.
- Little, Greg, Chilton, L.B., Goldman, Max, Miller, R.C. (2010). Turkit: human computation algorithms on mechanical turk. *ACM UIST*, 57-66