MobileWorks: A Crowd Platform That Works (Against Poverty)

Philipp Gutheim[1,2,3], Prayag Narula[1,2]
(with Anand Kulkarni[1,4] and David Rolnitzky[1])

ABSTRACT

Online labor marketplaces offer the potential to automate a variety of tasks too difficult for computers, but present requesters with significant difficulties in obtaining accurate results. We share experiences from building MobileWorks, a crowd platform that departs from the marketplace model to provide robust, high-quality results. Three architectural contributions yield measurably improved accuracy on input tasks. A *dynamic work routing system* identifies expertise in the crowd and ensures that all work posted into the system is completed with bounded completion times and at fair worker prices. A *peer management* system ensures that incorrect answers are prevented by experienced members of the crowd. Last, *social interaction techniques* give the best workers the ability and incentives to manage, teach and supervise other members of the crowd, as well as to clarify tasks. This process filters worker error and allows the crowd to collaboratively learn how to solve unfamiliar tasks.

INTRODUCTION

Human computation platforms are online services that allow software systems to subcontract portions of their functionality to a large crowd of human workers over the web. These systems have become increasingly popular in recent years, and new applications are discovered on an ongoing basis in problem domains ranging from vision and artificial intelligence to software development and business process outsourcing. Because they require software systems to interact with humans on a paid basis, human computation platforms provide an interesting set of challenges at the interface of economics, HCI, and theoretical computer science.

Human computation systems are typically assessed on their ability to provide accurate results. In human computation systems *accuracy* refers to the correct completion of tasks posted by a requester, by the requester's assessment of correctness. Quality control in terms of accuracy is a central challenge in all types of human computation platforms. Human computation systems may provide inaccurate results for one of several reasons. We classify some of these below.

1) Online workers may not want to do what they are being asked to do in exchange for the rewards offered, due to lack of motivation or malice. This is the *incentive* problem.

2) Even when workers are properly motivated to do a task correctly, the task may be ambiguous or not clarify what to do in all cases. This is the *task specification* problem.

3) Workers may make errors despite wanting to carry out a properly specified task. This is the *human error* problem.

The first two challenges are intertwined: when workers do not understand a task due to ambiguous design, they are likely to lose motivation.

Other failure modes originate from the internal structure of human computation platforms. For instance, many tasks posted to online marketplaces languish and are never solved because workers view their rewards as inadequate or cannot find the tasks in a marketplace, a problem termed starvation. These difficulties complicate efforts to incorporate human computation services into software, since it becomes difficult to provide bounded task response times and accuracy guarantees.

The most well-studied commercial crowdsourcing platform today, Amazon's Mechanical Turk, operates as a web-based marketplace where employers can post groups of tasks to be solved and workers can browse these tasks or choose to answer them. The problems of designing effective tasks, filtering unqualified workers, and eliminating incorrect answers are largely left to employers, with ratings available for filtering workers based on past accuracy. Most employers using Amazon Mechanical Turk must build extensive quality control infrastructure on top of narrowly-constrained application domains, such as audio transcription, web research, or text recognition, and employ domain-specific techniques to deal with possible errors.

---

[1] MobileWorks Inc. , 2140 Shattuck Avenue, Suite# 1110, Berkeley, CA

[2] School of Information, University of California, Berkeley

[3] Department of Computer Science, University of California, Berkeley

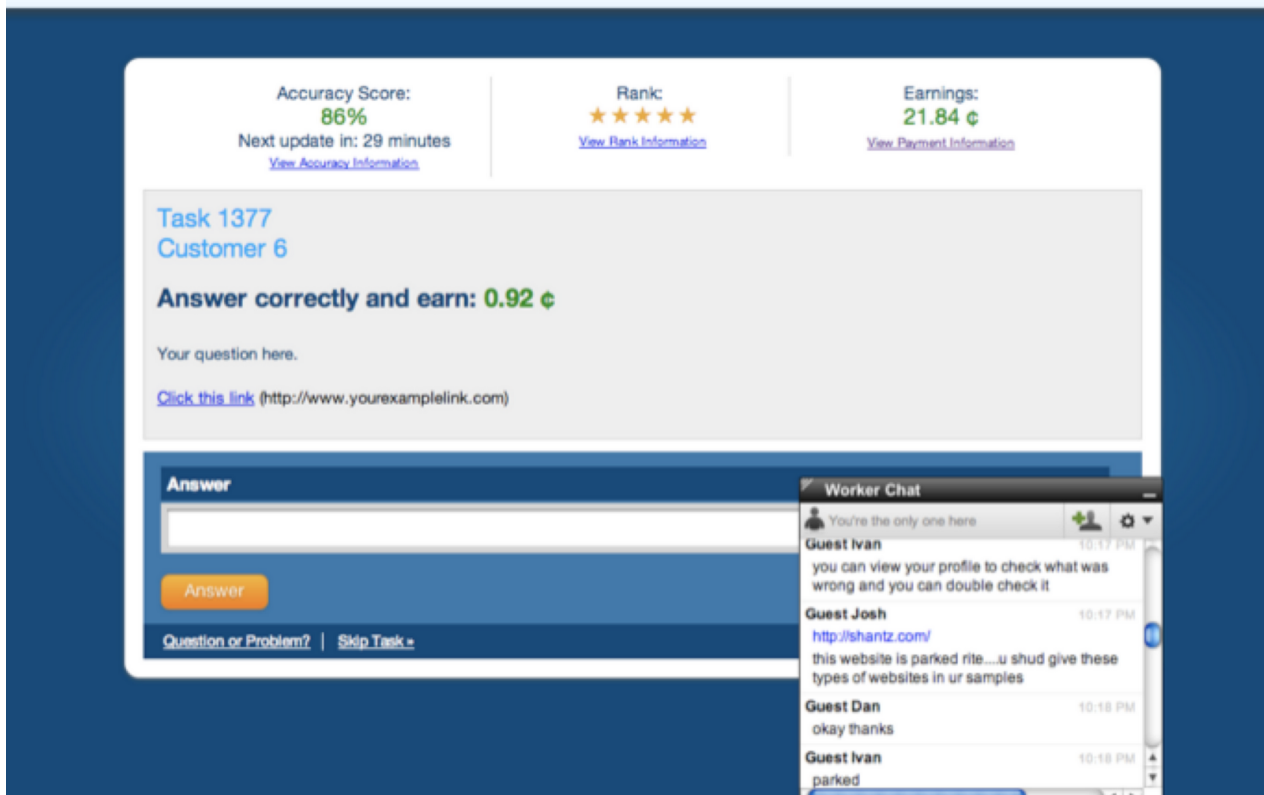[4] Department of IEOR, University of California, Berkeley

Figure 1. The MobileWorks interface serves a stream of appropriate tasks continuously to workers, interleaving qualification tests, expert tasks, unskilled tasks and training material. Workers are presented with immediate feedback on performance and earnings. Worker-to-worker chat is used to debug tasks on the fly, and workers can escalate tasks for managerial review.

MobileWorks is an alternative crowdsourcing platform designed to prevent the accuracy and speed deficiencies faced by employers in online labor marketplaces. A central difference is that MobileWorks is *not* a marketplace: it operates as an algorithmically managed service, routing work to qualified workers and recruiting additional participants as needed. Tasks presented to the system are automatically matched and presented to qualified workers identified through a combination of human and programmatic testing. Discrepancies are resolved by the best workers (managers) who play an active role in maintaining the quality of the system and managing other workers. Worker-to-worker interaction, led by managers, permits additional worker training and discussion of individual tasks. These mechanisms allow us to address the same class of tasks solved on conventional labor marketplaces, while providing substantially higher accuracy and shielding employers from the burdens of quality control.

MobileWorks operates with a social mission to provide employment to marginalized populations in the developing world; as such, the majority of its workers are drawn from low-income populations in South and Southeast Asia and paid a fair baseline wage. We discuss novel aspects of MobileWorks' architecture in several areas, including task routing, pricing, peer management, and worker-to-worker interaction. Following this discussion, we present a brief analysis of the 500,000 results produced in the lifetime of MobileWorks examining how effective these techniques were at maintaining quality in a real-world commercial task.

RELATED WORK

A substantial body of literature focuses on how to reduce worker error during human computation. In early work in human computation, the ESP Game [1] used redundancy to control error, finding that sending the same microtask to multiple workers could be effective in mitigating the effects of error and malice. Redundancy works well in mitigating human error with random structure, but is vulnerable to structural confusion in tasks and to worker collusion. Alternatives for maintaining quality include voting systems and iterative improvement [2]. Not all human computation platforms have taken the form of labor marketplaces. Alternatives include online games such as GWAP for labeling [3] and fold.it for protein-folding [4], as well as content-distribution networks that embed crowdsourcing tasks into other web activities such as CAPTCHAs or

sweepstakes [5]. While incentives are different, these systems face many of the same challenges in maintaining quality as marketplaces.

Other approaches such as Crowdflower's have layered redundancy with tracking of a worker's historical performance and ongoing worker assessment on so-called "gold standard" tasks with known answers [6]. Historical performance is readily gamed in open marketplaces and may not adequately predict worker performance on unfamiliar tasks [7]; gold standard tasks are unavailable for many kinds of work like content creation.

More sophisticated approaches use *workflows* that divide complex tasks such as essay-writing and editing into various short-duration tasks that are distributed among different workers. Soylent introduced the *find-fix-verify* workflow for complex text editing, supported by a peer review step that let workers verify that another worker's edits were carried out correctly [8]. Turkomatic introduced the price-divide-solve meta-workflow that uses workers to assist in designing workflows for complex tasks [9]. Soylent identified important gradients in quality and behavior among worker populations, and Turkomatic found that improper allocation of workers to tasks within a workflow could reduce quality of results. These results imply the possibility that filtering and matching based on categories of worker can yield enhanced accuracy.

Shepherd found that worker-to-worker feedback and requester-to-worker feedback displayed immediately after a task was useful in improving the quality of work on crowd platforms [10]; however, it did not investigate the impact of real-time worker interaction.

Several recent papers have explored mechanisms to reduce the latency of crowd algorithms, including combating task starvation and obtaining real-time results. VizWiz showed that streaming tasks to workers could maintain interest until relevant tasks become available [11], and found that using SEO-style techniques could enhance responsiveness of workers. Adrenaline showed that it was possible to pull results from crowd workers in seconds by keeping them "on call" until needed [12].
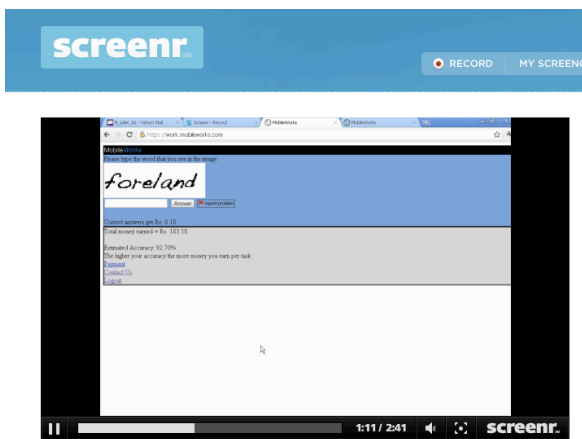

Figure 2. A training video produced independently by a manager in MobileWorks, discussing how to use the mobile-phone interface and how to deal with task discrepancies. These videos emerged organically as a consequence of our incentive structures.

The use of crowdsourcing with populations in the developing world has been explored before. TxtEagle, deployed in Kenya, used SMS text messages to provide tasks like audio transcription, local language translation and market research [13], but were limited by the capabilities of SMS to deliver more sophisticated tasks. SamaSource also seeks to employ marginalized populations with microwork, but manually optimizes outsourcing processes and partners with on-the-ground outsourcing agencies [14]. A 2010 study found that Mechanical Turk was largely unsuitable for workers at the bottom of the pyramid to find employment without additional modification [15] and caused additional issues of fairness [16, 17], even as other forms of online work such as e-lancing presented a rising opportunity for the developing world [18 19, 20].

System Architecture
We provide an overview of techniques used in the MobileWorks architecture. The central mechanism is an active *task routing system* that assigns tasks from a priority queue to individual workers in order to build consensus according to a standard quality assurance workflow. As part of this system, we have implemented techniques for identifying appropriate expertise within the crowd, for automatically pricing tasks, and for escalating errors or difficult instances to a special population of managers. Managers have high accuracy and proficiency, and we use them for worker recruitment, to evaluate potential problems with requester-defined tasks, and to resolve discrepancies in tasks.

Task routing and the dynamic work queue

Tasks are posted to MobileWorks via a REST API or through a web dashboard. At posting, requesters specify a set of instructions, a set of answer fields, preferences for what kind of workflow they require, and optionally, a set of skills. The set of skills is chosen from a list of keywords; alternately, requesters can define custom skills by providing gold standard tasks. Work pushed into MobileWorks is inserted into a priority queue of tasks that are processed by workers in turn.

When a worker arrives at MobileWorks via the web or a web-enabled mobile phone, he is assigned the next task in the queue whose skill requirements he meets and which that worker has not carried out yet. Workers whose overall accuracy is below a certain level are reassigned to training tasks until their accuracy improves.

Workers are assigned tasks, rather than selecting them from a list. We hypothesize that this interface enables workers to perform work more efficiently than a marketplace listing possible tasks, since they do not need to interrupt their work to search for subsequent tasks, and it provides a useful guarantee that every task will be answered. If tasks remain near the end of the priority queue for an excessive amount of time, we interleave them with tasks near the front by adjusting their priority to prevent large jobs from deadlocking the system. This means starvation is impossible so long as workers continue using the system.

Using a queue rather than a marketplace permits us to allow fine-grained control over the speed with which work is completed; by inserting work in the front of the queue, we can ensure that it is completed more quickly. If work is not completed due to the natural activity of workers in the system, emails are sent out to additional members of the worker population informing them that there are tasks available. While not as robust at scale as other approaches to control latency in real-time systems, providing a work queue affords a simple mechanism for preventing starvation.

Quality Control Workflows and Exception Handling

Since workers are assigned tasks rather than selecting them, it is crucial to build effective quality control mechanisms into the platform architecture. Each submitted task is inserted into the queue multiple times and presented to multiple workers via either a parallel or iterative model chosen by the requester when work is posted, depending on whether the task has an objective or subjective answer. In the iterative model, tasks are given to a sequence of workers in turn until a prespecified quantity of workers have sequentially edited the task. In the parallel model, if n distinct workers submit the same response to a given question, the answer is presumed to be correct and returned to the end user via a callback. If the workers disagree, the task is served to additional workers until a quorum is reached or an upper limit on the number of workers involved is hit, at which point the task is considered *ambiguous* and marked for review by managers.

Workers can elect not to do a particular task, but must provide a reason. Workers who are unfamiliar with a task or confused by its instructions can choose to report confusion via a "report problem" button on the interface, which requires workers to select a reason – whether the instructions were unclear, the instructions did not cover what to do in a particular instance, or whether individual links or resources were not available. This guarantees that all tasks in MobileWorks will terminate with some form of feedback to requesters rather than starving silently. This provides a kind of exception handling unique to MobileWorks: much as low-level errors in conventional programs can be caught by exception handlers, tasks that are not answered due to faulty design can be escalated to managers for potential resolution (Figure 3), and eventually back to the requester to help debug tasks.

Expert Sourcing and Qualification

A large number of tasks submitted to crowdsourcing systems require the attention of workers with specialized skills. Historically, expert workers have been identified either by automated tests or by constructing dedicated communities like 99Designs or StackOverflow that can attract these groups. MobileWorks supports this mode of operation, allowing requesters to specify qualification tests required for tasks. However, automatic test structures can be inadequate to identify optimal workers for complex skill categories like writing and content editing; for many open-ended tasks an automatic test cannot verify quality.

MobileWorks uses crowd-powered mechanisms to recruit additional expertise as needed. When a task requiring skills is posted in the system, the system checks whether enough expertise is available in the worker population to meet the demand. If not, tasks are posted asking workers to refer individuals in their network or from their own contacts who might be able to meet the requirements. A manager reviews a sample task (or full resume!) submitted by the referred worker to ensure they have the skills claimed.

Once at least one expert with a given skill is in the system, they are tasked with the problem of checking the work of other experts in a peer-review system. Would-be experts are given subjective tasks that can be assessed by other experts to determine eligibility. This strategy assumes that managers are capable of assessing quality for tasks they lack fluency in; a more comprehensive option that we may implement is to use requesters for this role.

Fixed Payment and Requester Pricing

We believe that one cause of worker malice and task starvation in crowd marketplaces is improper incentivization: a failure to be able to earn effective wages. A brief survey given by the authors to 100 workers on Mechanical Turk indicated that the most compelling complaint was low pay or failure to pay by customers. Part of the reason is that employers are often unaware of the appropriate market price of their task, and so price tasks at excessively low prices on existing marketplaces. Consequently, MobileWorks pays workers a wage per task that is set by the system, enabling them on average to earn fair or above-market hourly wages in their local zones.

The payout for a certain type of task is determined by dividing a target hourly wage by the average amount of observed time required to complete the task, excluding outliers. This pay structure incentivizes talented workers to work efficiently while ensuring average workers earn a fair wage. Worker payouts are set so that individual workers in India will earn an average above-market wage of INR 65 (roughly $1.50) / hour working on computers and INR 25 ($0.60) / hour working on mobile devices. Given that typical workers in MobileWorks earned less than $3/day prior to joining, this wage provides a compelling source of income. Gaming this system to reach inaccurate prices would require collusion between large numbers of workers. Requesters are charged a fixed price for each task based on a multiple of the worker wage required. If a requester's tasks start to vary significantly in time, the requester is notified and asked to approve the new price before work continues.

A few workers in the United States wrote to the authors requesting to join our system at the posted pay rates of $1.50-$2 / hour. This suggests that the payment model is attractive from a worker perspective even when pay is modest, though the authors believe high wages are important component in the design of a crowdsourcing platform. We expect to admit US-based workers at US wages in the near future. For tasks such as English video transcription where market prices can be as high as $1-2 per minute of audio, as well as tasks requiring high domain expertise or US citizenship, these wages should be well-matched by market rates for their services.

Figure 3. The interface presented to managers lets them determine appropriate errors on exceptional tasks to be returned to users, and to review performance of a team of workers they've recruited.

Workers are paid for each answer they provide that matches the answer determined by the overall system (through

redundancy and manager review); their payout is tiered, with workers whose accuracy is below 80% only earning 75% of their overall possible earnings. This was a choice to encourage long-term attention to accuracy rather than considering individual tasks – when a worker recognizes that his or her incorrect answers will affect his or her payout for correct tasks downstream, the cost of incorrect answers increases. Workers can view their list of incorrect responses via a profile page (Fig 4); if they object, they can report tasks for review by a manager, which further improves accuracy.

Social AND Management techniques

Social management techniques provide a new strategy for producing effective work in crowdsourcing platforms. These contributions are unique to MobileWorks' architecture. We facilitate worker-to-worker and manager-to-worker communication using tools within and outside the platform

Manager Recruitment and Supervision

MobileWorks grants managerial privileges to certain high-performing members of the crowd, termed *managers*. These workers play a central role in the efficient functioning of the architecture, as they supervise the work of other members and share in their earnings, providing an additional level of supervision and support. Managers in MobileWorks earn an additional 10-15% of what each worker they supervise earns, and are paid an additional per-task fee for each ambiguous task they resolve on behalf of the system (Figure 3).

This model initially emerged as a byproduct of our efforts to recruit overseas workers, but proved effective enough to use as a mechanism for maintaining quality. We identified the highest-performing members of our crowd by average accuracy and task volume and designated them as managers, asking them to recruit new workers. Peer recruitment allowed us to filter incoming workers based on demographic profiles and likely motivations; this process was largely controlled by managers, who were given a financial incentive to assemble effective teams. While the current crowd size of 500 workers is smaller than Turk's estimated 2000-10000 active workers, we expect that this kind of efficient filtration process can yield a comparably effective workforce over time.

Guest Ivan 10:17 PM
you can view your profile to check what was wrong and you can double check it

Guest Josh 10:17 PM
http://shantz.com/
this website is parked rite....u shud give these types of websites in ur samples

Guest Dan 10:18 PM
okay thanks

Guest Ivan 10:18 PM
parked

Guest Grace 6:40 PM
hi, if hoovers shows more than one with the same company name and there is no additional info in the instructions to crosscheck it with, like an address, how do i know which is it? ? AEE Solar, for instance.

Guest Grace 6:43 PM
this is guest grace still, what i did with the same situations before was to look up on google and use the website to find the correct info, but now, well i am wondering if i was correct when i did that coz i might be looking at the wrong website. help!
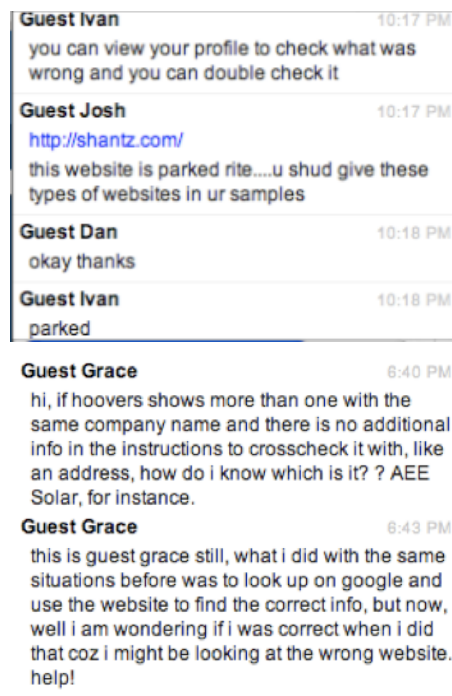
Figure 4. Worker-to-worker communication being used to clarify the user interface, the answer to a task, and the strategy used for solving it. Collaboration permits workers to debug tasks in situ and deal with unfamiliar tasks with managerial help without returning to requesters.

Managers proved effective in unexpected ways. At their own initiative, we observed that managers began using outside methods not facilitated by the platform, such as online screencasts and emails, to demonstrate to workers how to carry out work (Figure 2). These had the added benefit of institutionalizing knowledge about how to carry out work, which could be passed on to other workers. Because these methods seem to be effective, and because some workers have indicated they would use them more if embedded directly in the interface, we expect to add support for these interfaces in the future.

Worker-to-Worker communication

Worker error is often caused by an inability to understand the stated task or when instructions fail to cover corner cases that arise in a specific example of a task. To remedy this situation, we have implemented the ability for workers to communicate in real-time with a chat box embedded in the interface. Because our workers often reside in the same geographic areas and know one another in real life, this proved to be a useful tool to collaborate on work. Figure 4 shows instances where worker chat allowed workers to work around inadequate explanations in the task, to suggest additional examples that could be given to requesters, to teach other workers how to use the interface, and to confirm their theories about what a task meant.

SYSTEM PERFORMANCE

Aggregate Data

We have processed over 500,000 individual results since MobileWorks began field trials in March 2011, with the majority of tasks submitted by paid commercial customers and the remainder generated during testing and training periods. The posted tasks have varied from image tagging, processing, and speech recognition tasks, to web research and database population tasks, OCR and data extraction tasks, content creation tasks and simple planning tasks.

Few tasks (fewer than 5%) processed within MobileWorks relied on external quality assurance methods or workflows post-processing. While this is representative of the needs of customers using MobileWorks, it also indicates that the model succeeds in reducing the need for developers to engineer for quality on their end in many scenarios.

At its peak load, MobileWorks has produced roughly 20,000 answers per day. The average accuracy of individual workers in our system across all tasks is 85%, with the top quartile having an accuracy of 90-98%, the second quartile at 85-90%, and the third at 80-85%. To date, starvation has not occurred in MobileWorks, indicating that our architectural measures against it are effective – though we expect that our ability to qualify expert workers may bottleneck performance in the future.

Overall, the response of workers to MobileWorks has been overwhelmingly positive. We asked 30 web-based workers to rate the usability of the system on a five point Likert Scale. 30 out of 30 users rated the usability at a four or higher. Moreover, all users indicated that they would recommend the system to their friends and family, giving a 4 out of 5 or higher on a Likert scale. Workers said that the biggest advantage of the system was that the work could be done anywhere, at any time of the day. Typical comments include: "I could do the work and earn money while traveling to my regular job or even while watching television".

How Much does Managerial Review Help?

We empirically determined that the techniques managers apply in teaching workers play a useful role in training. However, managers play an essential algorithmic role in the quality assurance process, determining the final answer for tasks that are escalated due to being marked as ambiguous or unclear. Does this process provide a useful effect on quality? As a preliminary answer, we examined a set of 15,000 identical web research tasks for which the worker had to visit a website and search for a contact email address. If the link was broken or the website was parked, the worker had to mark the task as invalid. When two workers marked a task as invalid, it would escalate to a manager. If there was no email address available, the worker was to mark the task as 'information not available'.

Our results show that out of 800 escalated tasks, the managers found an email address for around 550 that would otherwise be sent back to the user, which is a substantial 70% improvement. In total, we did not find email addresses for around 5000 tasks (prior to the review). Using managerial review, we were able to find answers over 10% of these tasks. This indicates managers can prove effective in disambiguation.

MOBILEWORKS AS A FAIR-WAGE EMPLOYMENT PLATFORM

MobileWorks and mobile crowdsoucing

MobileWorks was designed to provide employement opportunities to the under-employed lierate population in underdeveloped areas around the world. MobileWorks was initially concieved as a mobile phone based crowdsourcing platform working on basic cell phones. [21]. In the cited study, we found that the participants increased their earning by 80-

100% . In addition, many participant were doing the tasks not on their mobile phones (as expected) but on their computers. When asked the workers mentioned that it was much faster for them to work on their computers (and hence much more profitable to them). On the other hand, it became increasingly hard to find meaningful tasks that could be done mobile phone. In an informal survey of the requesters we found that most of the task required access to a fully functional web-browser and a computer. Learning from the experience we developed a fully functional web interface for doing the tasks. The web interface has now served almost 95% of all tasks done by MobileWorks. [2]

The Social Mission
The mobileworks social mission aims at providing our workers above market wage for their tasks (www.mobileworks.com/fairtradework/). Since we are not a market place, the prices for doing a task are set by the system. This price is determined based on a fixed fair wage multiplied by the estimated time a worker takes to do the tasks. MobileWorks also allows workers to work whenever they want and from wherever they can access a computer. This has been mentioned by our workers as one of the main advantage of MobileWorks as a suplemental wage provider.

This phenomenon has been observed time and again by us in different studies. For example, we observed from our data that one of our members did not do any work on Thursdays while being regular on every other day (even weekends). On enquiry we found out that she was a house-wife and couldn't do any work on Thursday since it was the day of laundary in the house-hold. She also mentioned that she couldn't afford to do this if she took any other job elsewhere.

Crowdsourcing as a stepping stone to career path
To make MobileWorks as a fair employment platform, it is important for us to understand how our workers can use their MobileWorks experience to futher their professional and personal goals. We conducted study to find out what career path the crowdsourcing workers generaly pursued and if their crowdsourcing experience was already helping them in improving their skill set. We ran multiple studies (both quantitative and qualitative) on Amazon's Mechanical Turk. [3]

We found out that the aspirations of these workers were truely diverse. Though majority of the workers aspired to be better at crowdsourcing work, a lot of them wanted a career in IT, web development, education, finance and even health services.

Aspirations

| Choices	Percentage	Count
Mechnical Turk Work	27.78%	15
Computer or IT	24.07%	13
Own Business (Self Employed)	14.81%	8
Management Position	7.41%	4
Finance/Banking	7.41%	4
Content Creation	3.70%	2
Office Administration	3.70%	2
Other	11.11%	6
Business Process Outsourcing	3.70%	2
Student (Higher Education)	3.70%	2
Education (Teachers/Professors)	1.85%	1
Health Care (Hospitals etc.)	1.85%	1
Total		54
Unanswered		2

Figure 5. Results of a survey to find out the career aspirations of crowdsourcing workers. The results are diverse with answers such as Crowdsourcing work, Computer/IT, Self-employment and Banking.

However, when we tried to find out how these workers were trying to achieve these goals while doing the work, certain patterns emerged. We found out that most workers were trying to improve certain basic skill sets while doing the work. These skill-sets include, typing speed, english as a foreign languge and communication skills. Certain intersting edge cases such as creative writing also emerged during the qualitative study.

---

[2] This figure is skewed towards the web interface because most of the task that enter the system and marked to be not do-able on mobile phones and hence are not served through the mobile interface.

[3] The biggest crowdsourcing marketplace for micro-work on the web.

Skills that you learned from MTurk

Choices	Percentage		Count
Typing speed		75.00%	42
By making money		71.43%	40
English (or other language skills)		66.07%	37
Computer/Software Skills		64.29%	36
Cultural Knowledge about different countries(like the US)		50.00%	28
Communication Skills		44.64%	25
		Total Entries	**56**
		Unanswered	*1*

Figure 6. Results of a survey about the skills workers learned (opr preferred to learn) while doing crowdsourcing work. The options were chosen by analyzing responses of a qualitative survey. The qualitative survey repeatedly mentioned Typing Speed and English and repeatedly mentioned the fact that extra income helps in the career too.

In a qualitative study we also found out that some crowdsourcing task led to unexpected and sometimes tangentially related career opportunities to some workers. For example one of the participant in our study shared the following experience about his crowdsourcing work on mechanical turk.

"I'm currently getting my bachelor's degree in computer science and engineering. And we are now having on-campus recruitment from various firms. And my first interview was for the post of a creative writer for an IT firm. The interview board inquired me that whether I have any previous experiences. I told them I'm a Mechanical Turker at Amazon. I write articles so an so. They were quite satisfied and told me to send them the articles I have written for Amazon. I send them the articles and they were quite impressed with my work and appointed me as a creative writer in their editorial board."

Another crowdsourcing worker mentioned that he has been using his crowdsourcing experience to communicate and make friends with other workers from around the world.

"To communicate with other workers by accepting friend request like facebook."

Another result that emerged from the study was that there are a number of crowdsourcing workers who have specialized skills which they don't use in their crowdsourcing work. For example, a number of crowdsourcing workers were engineers, IT professional or even doctors by training.

Hence, crowdsourcing has been helping the workers futher their career in unexpected ways. With MobileWorks we will try to explicitly help workers in their career trajectory by providing them work which aligns with their current expectation from crowdsourcing tasks as a vehical to a career advancement. By helping our workers learn English as a second language, improve their typing speed and helping them develop soft skills like communication and better writing skills.

To futher the career of the more specialized members of the workforce it would be important to identify and utilize the skills that these workers already have. These skills can then be used in specialized crowdsourcing as mentioned in the section titled "Expert Sourcing and Qualifications". These skills can be furthur developed by providing these workers training and mentorship from other crowd members and managers.

Conclusion and future work

MobileWorks combines several known techniques for quality control with novel contributions in human computation in routing, pricing, management, and worker-to-worker interaction. Whereas the core ideas of routing, redundancy and review are not new to the crowdsourcing literature, their combination can provide an effective architectural alternative to the marketplace model.

Our experiences with MobileWorks suggest that as the crowd control systems place an increasing emphasis on accuracy, their architecture will begin to mirror the structure of traditional real-world firms, with collaboration and supervision playing a more important role in achieving accuracy than purely automated techniques. This is to be expected; crowd computing systems represent a mix of technological and organizational components. There are many additional improvements to the

system under development that mirror the processes used within organizations, such as introducing workflows and improving task standardization.

We have reported only a subset of results from MobileWorks to give a sense of the overall impact of the techniques we have used. However, we have a sizable body of data from over a half-million tasks executed in the system, as well as demographic studies carried out on the worker pool, which has expanded to nearly 500 workers in ten cities throughout the world. We look forward to making these available to the research community. MobileWorks is available for use online at www.mobileworks.com.

REFERENCES
1. Ahn, L. von, Dabbish, L. Labeling images with a computer game. CHI 2004.
2. Little, G. , et al. Exploring iterative and parallel human computation. HCOMP, 2010.
3. Ahn, L. von. Games with a Purpose. IEEE Computer 39, (2006), 92–94.
4. Cooper, S., et al. Predicting protein structures with an online game. Nature, 2010.
5. Ahn, L. von, et al. reCAPTCHA: Human-based Character Recognition via Web Security Measures.
6. Oleson, Sorokin, et al. Programmatic Gold: Targeted and Scalable Quality Assurance in Crowdsourcing. HCOMP 2011.
7. Ipeirotis, P. Mechanical Turk, Low Wages, and the Market for Lemons. 27 July 2010. www.behind-the-enemey-lines.com/2010/07/mechanical-turk-low-wages-and-market.html
8. 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.
9. Kulkarni, A., Can, M., Hartmann, B. (2012) Collaboratively crowdsourcing workflows with Turkomatic. CSCW 2012.
10. Dow, S., Kulkarni, A. Hartmann, B., Klemmer, S. (2012) Shepherding the Crowd Yields Better Work. CSCW 2012.
11. Bigham, J.P., Jayantc, C., Ji, H., et al. VizWiz: nearly real-time answers to visual questions. Proceedings of UIST 2010, ACM (2010), 333-342.
12. Bernstein, M., Brandt, J., Miller, R., Karger, D. Crowds in two seconds: enabling realtime crowd-powered interfaces. Proceedings of UIST 2011.
13. Eagle, N. (2009). txteagle : Mobile Crowdsourcing, Internationalization, Design and Global Development, 447-456.
14. Samasource. http://samasource.org
15. 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
16. Benjamin B. Bederson and Alexander J. Quinn. 2011. Web workers unite! addressing challenges of online laborers. In Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems (CHI EA '11). ACM, New York, NY, USA, 97-106.
17. Silberman, M.S., Irani, L. and Ross, R. 2010. Ethics and tactics of professional crowdwork. XRDS 17, 2 (December 2010), 39-43.
18. Lehdonvirta, V., Enrkvist M. (2011). Knowledge map of the virtual economy.
19. Ipeirotis, P. G. (2010). Analyzing the Mechanical Turk Marketplace. XRDS, 17(2), 16-21.
20. Ross, J., Irani, L., Silberman, M. S., Zaldivar, A., and Tomlinson, B. 2010. Who are the crowdworkers?: shifting demographics in Mechanical Turk. In Proceedings of CHI 2010, Atlanta GA, ACM, 2010.
21. Narula, P., Gutheim P., Rolnitzky D., Kulkarni, A., Hartmann B., MobileWorks: A mobileworks crowdsourcing platform for workers at the bottom of the pyramid. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011),* San Francisco, USA, 2011 (Poster)