# The Role of Super Agents in Mobile Crowdsourcing

**Mohamed Musthag** and **Deepak Ganesan**
Department of Computer Science
University of Massachusetts, Amherst
Massachusetts, U.S.A.

### Abstract

The ubiquity of smartphones have led to the emergence of mobile crowdsourcing markets, where several tens of thousands of smartphone users participate to perform tasks in the physical world. Mobile crowdsourcing platforms are uniquely different from their online counterparts in that they cater to smartphone users, require spatial mobility, and have tasks that involve more data collection and less human computation. Despite the emergence and importance of such mobile marketplaces, little to none is known about their dynamics.

This paper provides an in-depth exploration of several aspects of mobile crowdsourcing markets based on a year-long dataset from a leading mobile platform. We find that like online crowdsourcing markets, a small core group of workers account for a disproportionately large proportion of activity generated in the market. We find that these *super agents* are highly efficient, working on tasks more quickly and picking fewer lowly priced tasks. We discover that while all agents chain several tasks into one session, hence potentially amortizing travel costs, super agents are $3\times$ more likely to chain tasks. Unlike online crowdsourcing markets, we find a skew towards more males, an even younger population, and higher education levels.

## Introduction

The past decade has seen unprecedented growth in mobile phones, with millions of these devices becoming first-class citizens of the Internet. Phones have become increasingly sophisticated, with GPS and high-resolution cameras being widely available. These capabilities have enabled new paid mobile crowdsourcing applications, where individuals are paid to contribute data through their mobile phones as they move around in their day-to-day lives. Several instances of such mobile crowdsourcing systems have emerged commercially such as Gigwalk[1] and FieldAgent[2] — these systems pay users several dollars to do small tasks including photos of buildings or sites, price checks, product placement checks in stores, traffic checks, location-aware surveys, and so on. Given the increasing popularity of such marketplaces, their characteristics are important to understand.

While there has been significant prior work on online crowdsourcing markets such as the Amazon Mechanical Turk [3], mobile crowdsourcing markets have been little studied. They are uniquely different from their online counterparts. First, mobile crowdsourcing marketplaces cater to smartphone users rather than online users, hence it involves a different demographic and different incentive amounts compared to online systems. Second, they require mobility in the physical world as most mobile crowdsourcing tasks can only be completed at specific locations, hence location-dependent factors such as population density and access to public transportation are more important than broadband internet access. Third, the types of jobs in mobile crowdsourcing system are largely composed of data collection jobs (e.g. capturing images), and involve less human computation (e.g. labeling images). As a result, the characteristics of the mobile crowdsourcing market can be expected to be different in terms of the user demographic, role of incentives, and usage patterns.

Our goal in this paper is to provide an understanding of user behavior in *mobile* crowdsourcing markets. Specifically, our analysis is based on over a year of data from MobileCrowd[4], a popular mobile crowdsourcing system which has hundreds of thousands of tasks, and tens of thousands of agents.

While examining MobileCrowd user behavior, we observe a surprising trend — a relatively small core-group of users generate a disproportionately large fraction of the activity. Figure 1 shows that 10% of active agents are responsible for a remarkable 84% of total earnings, and for 80% of the total tasks done on MobileCrowd. This behavior is perhaps, not unlike other online crowdsourcing marketplaces: a study on Task.cn showed that 0.1% of active users were responsible for proposing 20% of the winning solutions (Yang, Adamic, and Ackerman 2008) (in comparison, 0.1% of agents were responsible for 10.78% of earnings, and 11.12% of completed tasks); studies of Amazon Mechanical Turk have shown similar behavior as well (Deneme 2009; Fort, Adda, and Cohen ; Ipeirotis 2010b; 2010a). The existence of such heavily skewed behavior on MobileCrowd



[1] http://www.gigwalk.com
[2] http://www.fieldagent.net
[3] http://www.mturk.com
[4] Name changed for anonymity

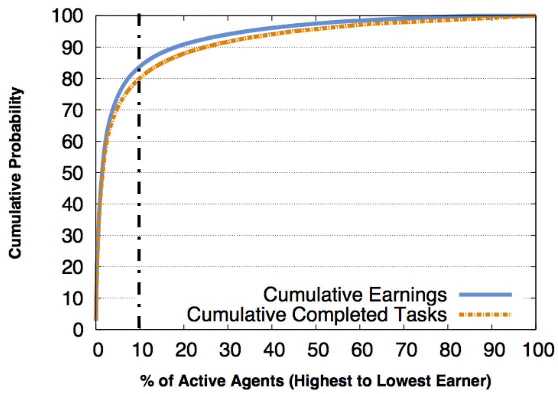


Figure 1: Cumulative Agent Earnings & Completed Assignments

makes it important to focus on this critical group of users. We define *super-agents* as the top 10% of active agents in terms of cumulative earnings.

Our analyses in this paper are largely shaped by this remarkable characteristic of the mobile agents on MobileCrowd. We examine several questions pertinent to user behavior through a super agent lens: 1) *Earnings* - Who are the most successful agents? How much do they earn? What techniques do they use to maximize earnings? 2) *Demographics* - Is this a young crowd looking for extra cash, working-class people looking for supplemental income, or retired folks looking to be engaged? Are these mostly well-to-do people doing it for fun or does the money matter? Are the super agents significantly different in their demographic make-up? 3) *Time spent on MobileCrowd* - Are these weekend warriors, people who work during workdays after-hours, or people who work primarily as agents? Do the super agents spend more or less time doing tasks? What about searching for tasks? 4) *Task popularity* - What types of tasks are popular among agents? Do super agents prefer a particular type of task? 5) *Location characteristics* - where do super agents typically do tasks compared to the overall population? 5) *Retention behavior* - Do super agents participate regularly in the market, or do they participate in bursts?

We find that super agents are highly efficient, working on tasks more quickly and picking fewer lowly priced tasks. We discovered that while all agents chain several tasks into one session, hence possibly amortizing travel costs, super agents are 3× more likely to chain tasks. Unlike online crowdsourcing markets, we find a skew towards more males, an even younger population, and higher education levels.

## Dataset

MobileCrowd is a smartphone-based crowdsourcing platform. It allows individuals with smartphones easy access to flexible work, while allowing businesses to tap into a mobile workforce. MobileCrowd is active in several major metros in the US including Los Angeles, New York City, San Francisco Bay Area, Seattle, among others. MobileCrowd customers post tasks, which usually have more than one instantiation and may be assigned in various locations. For example, a task could be to try a product and report on the customer experience and this task may be assigned for multiple stores in different locations. One of the biggest categories of tasks in the dataset is the Photosynth category. It requires agents to capture high quality panoramic images using the Microsoft Photosynth[5] app. Cafe and restaurant tasks have agents visit cafes or restaurants, take pictures of their menus, and may involve reviewing the quality of service, food, and atmosphere. Traffic tasks may have agents investigate the existence of a roadblock at a particular location.

Workers can search for tasks easily on a map, clicking through to view details of the task. Almost all tasks (¿99.8%) required agents to be present at the location of the task before they can accept one. Upon downloading the app, agents report their gender, age, highest level of education, and current profession, providing valuable information about the demographics of the agents on MobileCrowd. All activity generated by the agent on the app is logged with a timestamp and location. This includes agent authentication events, details of when they downloaded or viewed tasks, and information about when they started/completed a task. This information allows us to classify when and where agents were using the app, looking for assignments or actively doing work and forms the basis of much of our analysis.

MobileCrowd was initially released on the Apple iPhone, and has subsequently been released for the Android OS. We have data for over 400 days of MobileCrowd agent activity ending in September 2011. Our trace ended before the Android app was released so the dataset contains only iPhone users. Hundreds of thousands of different assignments were posted over this period, out of which tens of thousands had been completed by the end of our trace. Out of the over tens of thousands of unique agents who were registered to use the service, several thousand agents successfully completed at least one assignment, whom we call *active* agents. All analyses presented in the paper are over these active agents, and a subset of these active agents whom we refer to as super-agents, unless mentioned otherwise.

## Results

We analyzed agent activity with respect to the goals of the paper as discussed in the preceding section. We now present the results of our findings and their interpretations. We use mixed-effects regressions analyses, which account for correlation of observations within individuals, to test for significant differences. Where space allows, p-values are reported with the means and confidence bounds of the corresponding distributions.

### Earnings

First, we look at agent earnings and delve deeper to try and understand what, specifically, about super agents' *behaviors* distinguishes them from the rest of the crowd in terms of their earning potential. Is the higher earning simply an artifact of doing more tasks and spending more time doing

[5]http://itunes.apple.com/us/app/photosynth/id430065256?mt=8

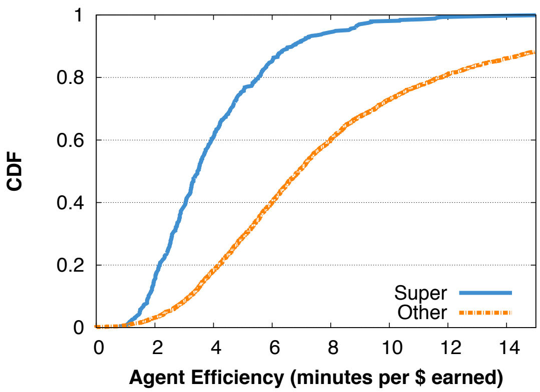


Figure 2: Working Efficiency

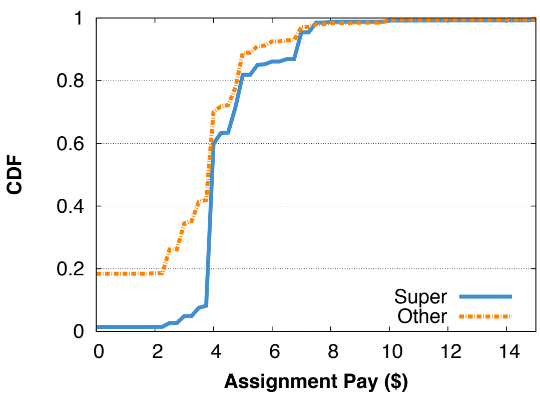


Figure 3: Completed Assignment Pay CDF

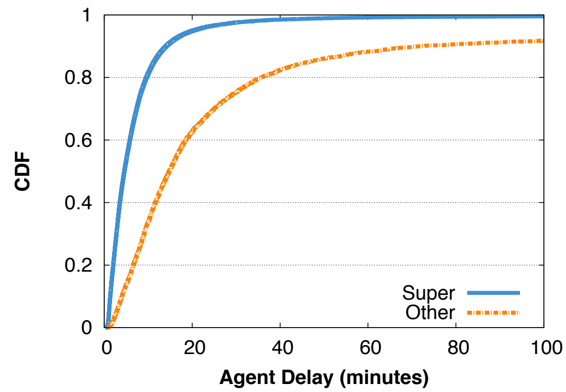


Figure 4: Completion Delay CDF - Photosynth Assignments

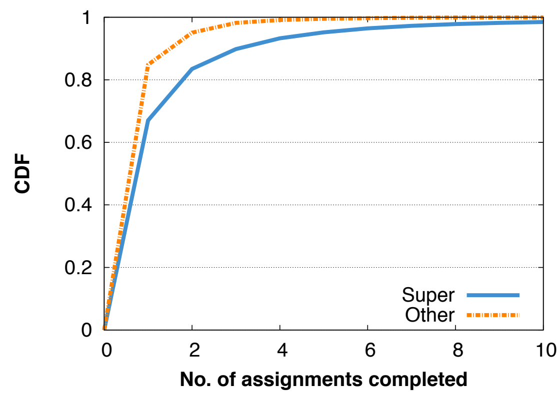


Figure 5: Session Worked Count CDF

tasks? Or are they actually more efficient in either searching for assignments, and/or working on assignments?

**Sessions** In order to answer these questions, we first break out the agent activity logs into sessions of activity. A session is defined as a series of agent activity events whose inter-event gap is not longer than 10 minutes. We chose 10 minutes as the threshold after examining many agent traces. Finally, we compute statistics about each session including duration, earnings, no. of assignments completed, and distance covered.

**Working efficiency** In order to measure how well an agent uses her time while working, we define *working efficiency* as the number of minutes she spends working on assignments in order to earn 1 dollar. Note that the inverse is the classical wage of the agent. Figure 2 shows the CDF of working efficiency for all active agents. Super agents spend on average $3.45$ minutes to earn $1$ dollar ($17.39/hr), while other agents spend $6.92$ minutes to earn the same dollar ($8.67/hr) $(p < 0.000)$

This indicates that super agents are more efficient at working on assignments, but we have yet to explain what leads to the higher efficiency we observe. Is it because a) super agents choose to do tasks that offer larger rewards? b) super agents are simply faster at doing these tasks? c) super agents discovered a different recipe besides a) and b).

Figure 3 shows the CDF of offered rewards on completed assignments. We see that super agents do in fact work on fewer lowly priced tasks ($p < 0.000$).

In order to see if super agents are faster at doing tasks, we look at the largest category of tasks in MobileCrowd, which are the photosynth tasks and check if super agents are quicker to complete these tasks. We fix the category of tasks to control for any differences between tasks in terms of the amount of work involved. Figure 4 shows the CDF of time taken by agents to complete these photosynth assignments. While the rest of the crowd takes $14.75$ minutes on average to complete these photosynth tasks, super agents take merely $4.58$ minutes on average to complete these tasks $(p < 0.001)$. Super agents are more than 3 times more efficient at doing photosynth tasks than the rest of the crowd, contributing to their remarkably high working efficiency.

**Chaining Assignments** One of the distinguishing characteristics of mobile crowdsourcing is that agents must be mobile in order to participate in the market, thus incurring an additional cost in travel time. One question then is whether

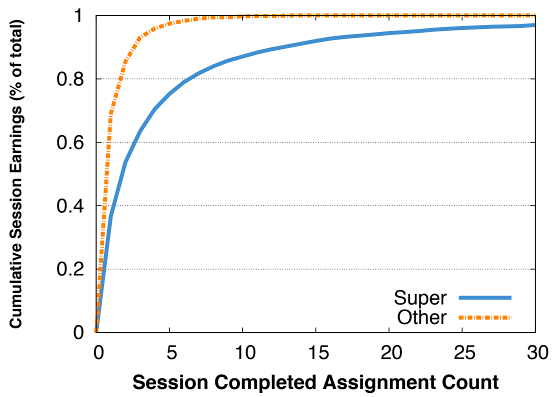


Figure 6: Cumulative Session Earnings

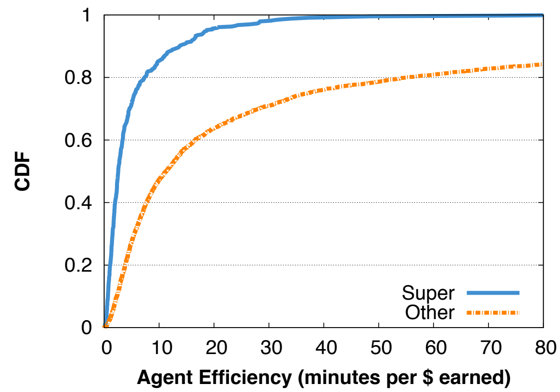


Figure 7: Searching Efficiency

agents try to amortize this associated travel cost. Do agents choose to travel to areas where there are clusters of tasks available and chain-work on several of them in one trip?

Figure 5 shows that while non-super agents chained more than one assignment in only 5% of worked sessions, super agents chained more than one assignment in almost 20% of worked sessions - a four-fold difference. Indeed, we find super agents are $3\times$ more likely to chain multiple assignments into one ($p < 0.000$).

This difference is even more significant if we look at the proportion of total earnings generated by sessions where chaining occurred. Figure 6 shows the cumulative fraction of the agents' earnings produced by sessions where they completed one assignment, two assignments, and so on. For super agents, this plot shows that this 20% of worked sessions, where more than one assignment was completed, is responsible for nearly 50% of total super-agent earnings. For the rest of the agents, this plot shows that the corresponding 5% of sessions generated nearly 15% of non-super-agent earnings.

Thus, we see that chaining of work is indeed present in the MobileCrowd agents and that the most successful agents chain more assignments together in order to achieve higher working efficiency.

Note that to begin almost all tasks completed on MobileCrowd (99.34%), agents had to be present at particular locations. Therefore it is highly unlikely that agents could appear to chain tasks simply by completing tasks that did not require a geo-lock. Furthermore, as described earlier, we used 10 minutes of inactivity to demarcate session boundaries. We think that 10 minutes is not long enough for an agent to move a significant distance. Thus, we think it is very likely that completed tasks within sessions are spatially nearby. However, to make a definitive conclusion, we need to correlate the geo-locations of completed tasks within sessions.

**Searching efficiency** The time agents spend working on assignments is just one aspect of agent efficiency. The other aspect of efficiency is the time agents spend planning and looking for attractive tasks on the application. In this section we look at these sessions in which agents do no work, but spend searching and planning. Similar to *working efficiency*, we define *searching efficiency* as the amount of time an agent spends planning and searching for tasks to earn 1 dollar. Figure 7 shows the CDF of *searching efficiency* of agents. We find that super agents spend 2.6 minutes on average planning and searching in order to earn 1 dollar, whereas the rest of the agents spend 11.5 minutes on average planning and searching to earn the same 1 dollar. ($p < 0.000$)

Thus, not only are super agents choosing to do higher paid tasks, working more quickly, and chaining assignments to amortize travel costs, they are also planning and searching for tasks more efficiently.

These efficiency results do not take in to account the duration of time agents have been using the application. Consequently, this could be the result of super agents learning to become more efficient over time. On the contrary, however, we do not find any evidence of this in the data. Therefore, super agents' higher participation should not have any more positive effect on efficiency than for the rest of the agents.

## Time of use

Next, we try to answer questions related to agents' usage behavior in terms of when they perform tasks. How much work gets done during the weekend vs. during the week? Do agents work after their regular jobs on weekdays? Are the super agents' usage patterns different from the rest of the agents?

**Day of week** Figure 8 and figure 9 show histograms of worked sessions by the day of the week for super agents and non-super agents respectively. We find no evidence in the data to suggest that super agents are different than the rest in terms of what day of the week they choose to do work. For both groups, weekdays see more activity than weekends.

**Hour of day** Having discovered that lots of worked sessions fall during the week, we turn to see at what times during the day work gets done. Figure 10 shows the histogram of the hour of day when work was done over all agents during the weekday. We see that most of the work gets done in the afternoon with activity starting to pick up at noon and peaking at 5pm. This suggests the presence of two groups of

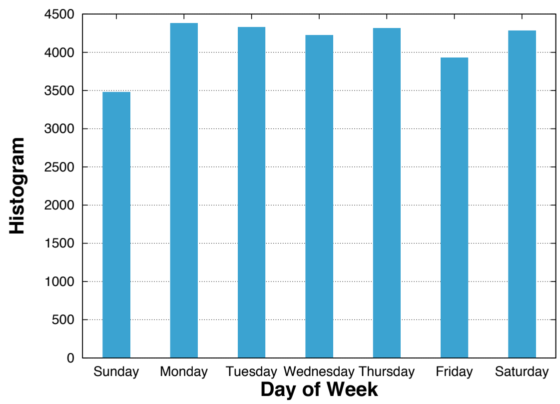


Figure 8: Histogram of time of use - Day of week - Super agents

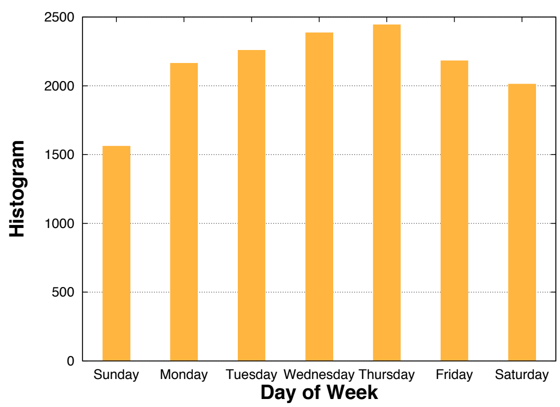


Figure 9: Histogram of time of use - Day of week - Other agents

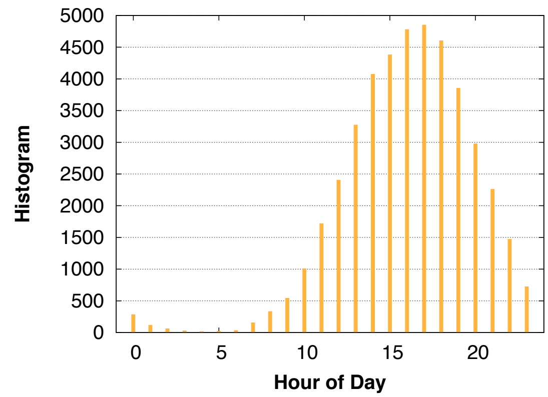


Figure 10: Worked Sessions by Hour of Day - All agents

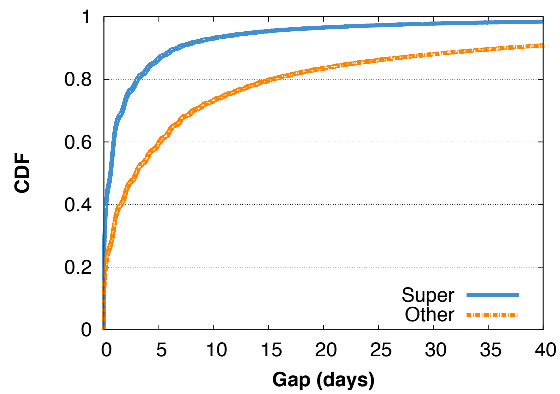


Figure 11: Inter-Worked Session Gap CDF

people: 1) those who are mostly free-lancers and have flexible work schedules; 2) those who work on MobileCrowd after regular working hours. From this aggregate data, we cannot tell what fraction either of these two groups command. We do not find any difference between the super agents and the rest of the crowd in terms of the hour of day work gets completed during the weekend or weekday.

## Frequency of use

Now we look at how frequently agents do work on MobileCrowd. How does interest in the mobile marketplace change over time? What is the mean time between worked sessions? Do super agents work more or less frequently than non-super agents? Figure 11 shows the CDF of the gap between two worked sessions from all those agents who have completed at least 2 worked sessions. We find that the mean time between worked sessions is 12.67 days for non-super agents, and the corresponding figure is 4.37 days for super agents on MobileCrowd ($p < 0.000$), suggesting that super agents are more engaged and use MobileCrowd to complete tasks more frequently. Figure 12 plots the time assignments were completed on a timeline for 3 randomly picked super agents (top 3) and 2 randomly picked non-super agents (bottom 2) in NYC. This figure also shows that non-super agent working activity is more diffused than super agent working activity, which tends to exhibit more bursty behavior.

## Demographics

We now report on the demographics of the active agent pool. Table 1 show that there are more men (71%) than women (29%) on MobileCrowd. Although slightly more men become super agents, the difference is not statistically significant.

Table 2 shows that agents on MobileCrowd are highly educated. Over 75% of agents hold a college degree, and over 20% hold an advanced degree. We find that more educated agents are significantly more likely to become super agents ($p < 0.010$).

Table 3 shows that most agents are young. Almost 70% of active agents are under 35 years of age. We find that older agents are more likely to become super agents ($p < 0.000$).

Table 4 shows two groups of agents: students (16.32% of active agents), who made up the largest profession on MobileCrowd, and photographers (3.75% of active agents),

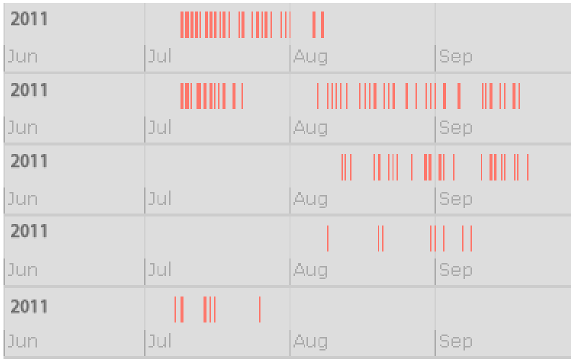


Figure 12: Completed Assignments Over Time. Top 3 slates show 3 randomly selected super agents; bottom 2 show 2 randomly selected non-super agents

| Gender | Super | Other | Total |
|---|---|---|---|
| **Male** | 74.31% | 70.58% | 70.95% |
| **Female** | 25.69% | 29.42% | 29.05% |

Table 1: Proportion of active agents by gender

who had the highest completed assignment yield, i.e. average number of assignments completed per agent. Despite being the largest profession on MobileCrowd, students are less likely to become super agents ($p < 0.009$).

Photographers had a completed assignment yield of 23.57 assignments, while the overall active agent population had a yield of 6.93 assignments. Indeed, we find photographers to be more likely to become super agent ($p < 0.036$). This behavior may be attributed to the high availability of assignments that involve capturing images (83% of all tasks).

## Online vs. Mobile Crowdsouring Markets

Worker behavior on online crowdsourcing systems have been extensively studied in the literature (Adda and Mariani 2010; Brabham 2008; Deneme 2009; Ipeirotis 2010b; 2010a; Mason and Watts 2010; Ross et al. 2010; Yang, Adamic, and Ackerman 2008). To the best of our knowledge, our findings represent the first view into the demographics and user behavior on mobile crowdsourcing platforms.

### Demographics

While a few of the demographic make up and behavioral patterns we observe on MobileCrowd are consistent with their online counterparts, many are different. Whereas women (65%) outnumbered men (35%) amongst *MTurkers* in the U.S. (Ipeirotis 2010b), we find the opposite on MobileCrowd, there are signficantly more male (71%) agents than female (29%) agents. Perhaps this difference is due to the high mobility requirement in mobile crowdsourcing markets. Yet, we did not find evidence that suggested that active male agents had any higher probability of success on MobileCrowd.

| Education | Super | Other | Total |
|---|---|---|---|
| **Some High School** | 0.93% | 1.02% | 1.02% |
| **High School** | 3.70% | 5.79% | 5.58% |
| **Some College** | 14.81% | 17.52% | 17.25% |
| **College** | 55.79% | 54.36% | 54.5% |
| **Some Graduate School** | 4.86% | 4.99% | 4.98% |
| **Graduate School** | 19.91% | 16.31% | 16.67% |

Table 2: Proportion of active agents by highest education level

| Age | Super | Other | Total |
|---|---|---|---|
| **Under 18** | 1.16% | 1.21% | 1.20% |
| **18-35** | 57.64% | 68.66% | 67.56% |
| **35-50** | 34.26% | 24.98% | 25.91% |
| **50-65** | 6.48% | 4.84% | 5% |
| **Over 65** | 0.46% | 0.26% | 0.28% |

Table 3: Proportion of active agents by age-group

(Ipeirotis 2010b) found that 54% of MTurkers in the U.S. to be between 21-35 years old. In MobileCrowd, the corresponding figure is over 62%, suggesting that whilst crowdsourcing in general is most popular amongst the young, *mobile* crowdsourcing is slightly more so. Despite this skew, and somewhat surprisingly, we find that the older agents have a higher probability of success on MobileCrowd.

Whereas (Ipeirotis 2010b) found 55% of MTurkers in the U.S. reported they hold at least a college degree, in MobileCrowd, over 75% reported holding at least a college degree, indicating a significantly more educated labor market in the mobile crowdsourcing sphere. These last two findings may be the result of the higher barrier to entry into the mobile crowdsourcing market. Many have access to the web, but not as many have access to smartphones.

HIT (Human Intelligence Tasks) completion activity on MTurk was less affected by weekends (Ipeirotis 2010a), which we find is not generally consistent with agent activity on MobileCrowd, but consistent with the behavior we observe on the part of super agents. This could be explained by the increasing fraction of MTurkers who view turking as a primary source of income, similar to the behavior we see from super agents who spend a significant portion of the working day working on MobileCrowd.

### Super Agents

A heavy tail of participants who have a significantly lower level of activity compared to the top contributors is not uncommon for any online community (Ipeirotis 2010a). In *online* crowdsourcing, MTurk, one of the most well-studied platforms, was found to exhibit this heavy-tail characteristic (Ipeirotis 2010a). In one longitudinal study (Ipeirotis 2010b), 10% of the most active MTurkers were found to complete 75% of the HITS. (In MobileCrowd, 10% of the most active agents completed 80% of all completed tasks.) The same study reported that 16% of the most active MTurkers earned 70% of total income. (In MobileCrowd, 16% of

| Occupation | Super | Other | Total |
|---|---|---|---|
| **Student** | 11.81% | 16.83% | 16.32% |
| **Photographer** | 5.79% | 3.52% | 3.75% |

Table 4: Proportion of active agents who are students and photographers

the most active agents accounted for 89% of total earnings.) This characteristic has been independently verified in (Deneme 2009), which found that the top 22% of MTurkers on AMT completed 80% of the HITs. Such a long-tail phenomenon is not only exhibited by the MTurkers (Ipeirotis 2010b; Yang, Adamic, and Ackerman 2008), but also seen on the part of *requesters* (Ipeirotis 2010a). Indeed, it was observed by (Ipeirotis 2010a) that 0.1% of total requesters on AMT accounted for more than 30% of the overall tasks posted on the market.

Yang et. al found a similar effect on a popular WitKey website – Task.cn. A WitKey website is an online knowledge market where users post questions or problems and other users provide solutions competing for a monetary award. Task.cn is one of the biggest WitKey websites in China with millions of registered users. They found that 0.1% of active users were responsible for proposing 20% of the winning solutions(Yang, Adamic, and Ackerman 2008) (by comparison, 0.1% of workers were responsible for 10.78% of earnings, and 11.12% of completed tasks) This 0.1% of users were additionally found to increase their win to submission ratio over time.

## Conclusions & Future Work

In this paper, we analyze user behavior on a popular commercially available mobile crowdsourcing platforms. As far as we know, our findings represent the first view into user demographics and behavior patterns in *mobile* crowdsourcing markets. Not unlike online crowdsourcing platforms, MobileCrowd also exhibited a super-agent phenomenon where just 10% of the top agents accounted for 80% of all completed assignments and 84% of all earned income. We discovered that these super agents were able to achieve such remarkable performance because they were more efficient using their time, not just working on tasks, but also possibly planning and searching for tasks. Additionally, super agents completed fewer lowly priced assignments and completed tasks more quickly than the rest of the agents. Because almost all assignments on *mobile* crowdsourcing markets require users to be present at particular locations, there is value in chaining tasks into one trip. We find chaining behavior to be present in all agents to some degree, but that super agents are $3\times$ more likely to chain assignments in to one session, hence possibly amortizing the travel cost across several tasks.

Unlike MTurk, which is an online crowdsourcing marketplace, we find that there are more men present in MobileCrowd, possibly due to the mobile nature of mobile crowdsourcing. In addition, we find that the MobileCrowd population is even younger than the MTurk population, and even more highly educated. These two factors may be attributed to the relatively higher barrier to entry in to the mobile crowdsourcing space. Many have access to the web, but not as many have access to smartphones. Both, freelancers with flexible work-schedules and people who have steady employed are present in MobileCrowd. We find that super agents are more engaged and complete tasks more frequently than the rest of the crowd, with a mean time between worked sessions of 4.37 days – almost three times more frequent than other active agents.

We have merely scratched the surface of understanding mobile crowdsourcing markets. Many compelling questions remain in this space, and we plan to tackle them in future work. One direction is modeling delay for tasks i.e. the time between task posting to completion given the location of the assignment and offered reward, based on super-agent and task availability in the geographic vicinity. Another direction is exploring the dynamics on the demand side of the market. What is the arrival rate of new tasks in the market? Is there a stable flow? Do these task posters achieve satisfactory results from the agents? And so on.

## References


Adda, G., and Mariani, J. 2010. Language resources and amazon mechanical turk: Legal, ethical and other issues. In *Legal Issues for Sharing Language Resources workshop*. Malta: LREC 2010.

Brabham, D. 2008. Moving the crowd at istockphoto: The composition of the crowd and motivations for participation in a crowdsourcing application. First Monday 13, 6.

Deneme. 2009. How many turkers are there? In *Deneme Blog*. http://groups.csail.mit.edu/uid/deneme/?p=502.

Fort, K.; Adda, G.; and Cohen, K. B. Amazon mechanical turk: Gold mine or coal mine? volume 37, 413–420. Cambridge, MA, USA: MIT Press.

Ipeirotis, P. G. 2010a. Analyzing the amazon mechanical turk marketplace. volume 17, 16–21. New York, NY, USA: ACM.

Ipeirotis, P. G. 2010b. Demographics of mechanical turk. New York, NY, USA: New York University, Stern School of Business.

Mason, W., and Watts, D. J. 2010. Financial incentives and the ”performance of crowds”. volume 11, 100–108. New York, NY, USA: ACM.

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 the 28th of the international conference extended abstracts on Human factors in computing systems*, CHI EA ’10, 2863–2872. New York, NY, USA: ACM.

Yang, J.; Adamic, L. A.; and Ackerman, M. S. 2008. Crowdsourcing and knowledge sharing: strategic user behavior on taskcn. In *Proceedings of the 9th ACM conference on Electronic commerce*, EC ’08, 246–255. New York, NY, USA: ACM.