Can Mobile Workforce Revolutionize Country-Scale Crowdsourcing?

Utku Günay Acer, Marc van den Broeck, Marc Godon, Claudio Forlivesi and Fahim Kawsar
Nokia Bell Labs
firstname.lastname@nokia-bell-labs.com

ABSTRACT

Traditional urban-scale crowdsourcing approaches suffer from three caveats - lack of complete spatiotemporal coverage, lack of accurate information and lack of sustained engagement of crowd workers. We argue that mobile workforces roaming around the city (and the larger country) can overcome all three caveats if their daily activity routines embed crowdsourcing tasks. To this end, in this paper, we report a first-of-its-kind study in which we explore behavioral attributes of mobile postal workers both quantitatively (6.3K) and qualitatively (6) to assess the opportunities of leveraging them for country-scale crowdsourcing tasks. Based on our observations, we develop a crowdsourcing infrastructure with carefully designed data collection strategies, and a corresponding wearable data collection application. We briefly present this solution and discuss its potential in country-scale crowdsourcing applications.

1. INTRODUCTION

With the emergence of Internet of Things (IoT), city governments increasingly use connected sensors to understand cities to enable planning for the future societal, economic and environmental wellbeing of their citizens. There have been a number of deployments of diverse IoT systems in urban spaces that offer a quantitative view of the urban landscape (e.g., Noise, Air Pollution, Mobility, etc.). However, such systems fail to capture the qualitative aspect of the urban landscape including the relationship between the citizens and their cities.

In order to get a better understanding of their cities, municipalities encourage their citizens to use crowdsourcing tools to provide information about the areas they live and support a more sustainable city. Through crowdsourcing, user-generated content can be collected and analyzed to identify and address problems cities face. The rise of location-based services and the widespread use of mobile devices makes such citizen participation more feasible and accessible [2]. In addition, companies such as Gigwalk[1], FieldAgent[2] and TaskRabbit[3] pay users to perform tasks such as price checks, audits, etc.

Mobile crowdsourcing attract considerable attention from the literature as well. It has been shown that crowdsourcing workers prefer to perform tasks outside the business hours when they have free time [1]. In terms of spatial scope, poorer neighborhoods may be avoided by workers [5]. Worker performance can be enhanced by bundling tasks together instead of assigning singular tasks to them [3].

Traditional approaches to urban scale mobile crowdsourcing suffer three major problems. Because participation is voluntary, the citizens mostly contribute information at times/ locations that are convenient to them. This leads to the *lack of spatiotemporal coverage* where most of the data is collected from highly populated areas like city centers at busy hours. The volume of data associated with suburban areas and off-hours is, on the other hand, is much less. Another issue of mobile crowdsourcing is the *lack of accuracy* in the collected information since there is no way of inspecting the users' input. As a result, these systems cannot guarantee the quality of the data. Finally, such systems have to deal the *lack of sustained engagement* where workers exhibit early enthusiasm but lose interest and drop out in time. As a result, only a small number of people introduce the large portion of collected data.

We argue that mobile workforces that roam around the city (and the country in the larger scale) can overcome all these three issues if it is possible to embed crowdsourcing tasks to the daily routines of the worker. In this paper, we focus on Belgium postal workers that collectively traverse an entire country to deliver letters and parcels. We present a study that quantitatively and qualitatively explores the behavioral attributes of the postal workforce. It includes an analysis of a dataset that reports information from more than 6300 postal workers delivering parcels to the households across Belgium. In addition, we present an observational study of six workers to assess opportunities in their behavior for crowdsourcing tasks. Our study suggests that postal workers can engage with crowdsourcing tasks without affecting their primary tasks.

Based on our observations, we design a crowdsourcing infrastructure that is tailored to the behavior of the workers. The system includes a wearable data collection application for the worker and carefully designed data collection strategies from this application. We briefly present this platform and discuss its potential.

2. BEHAVIOURAL STUDY

In this section, we present a study, both qualitative and quantitative, on the feasibility of leveraging a mobile workforce for crowdsourcing. Our use case in this study is the postal workers of Belgian Post Group (bpost).

[1] `http://www.gigwalk.com`
[2] `https://www.fieldagent.net`
[3] `https://www.taskrabbit.com`

WPA'17, June 19, 2017, Niagara Falls, NY, USA

DOI: `http://dx.doi.org/10.1145/3092305.3092309`

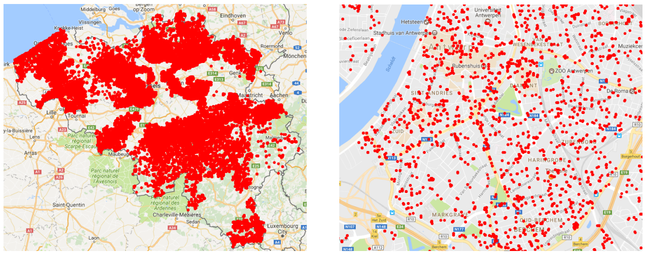

Figure 1: Delivery locations on a single day across Belgium and downtown Antwerp

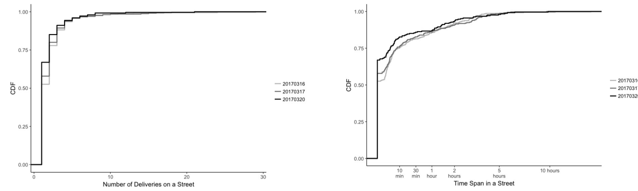

Figure 2: Distribution of number of deliveries and time difference between first and last delivery in each street. Most streets see only a small number of deliveries with small duration

bpost employs the largest payroll in Belgium including postal workers that deliver letters, packets and parcels to the 4.7 million households across the country. These workers go on 9200 mail rounds every day, delivering up to 9 million items[4].

2.1 Methodology

Our analysis consists of two elements. First, we analyze a dataset provided by bpost collected from a subset of workers who are equipped with mobile devices. The dataset yields timestamped information about *events* including barcode scans for delivered parcels to give feed back to senders and addressees. Then, we present a observation study on worker workflow and behavior.

We seek to understand whether the behavior of the postal workers during their rounds allow them to engage with a mobile device that they can use to provide responses to the queries about their current location. Based on this analysis, we make design decisions about the information collection system. Specifically, we address how and when the queries are best presented to the workers.

2.2 Quantitative Analysis

In this section, we present an analysis of a dataset provided by bpost. This dataset included a list of *events* from mobile terminals given to a subset of workers who deliver large packets and parcels to citizens. These events include scanning of barcodes on delivered items as well as other supplementary information such as charging of the device, etc. Events may also be geotagged with GPS coordinates.

The traces provide information about more than 6300 workers[5] with around 200K events in each of the three days. Approximately half of the events correspond to deliveries. However, not every event is tagged with GPS coordinates. In average 55000 events per day include GPS coordinates and about 40000 of them correspond to deliveries.

As shown in Figure 1 (a), the deliveries with GPS coordinates are not uniformly distributed across the county. bpost divides Belgium geographically into 242 distribution centers. The deliveries from some mail centers have no valid GPS coordinates. On the other hand, most deliveries from some mail centers, such as Antwerp Center, have valid GPS coordinates as shown in Figure 1 (b). During the course of these three days, 93% of deliveries associated with this mail center have valid GPS coordinates. For convenience, our analysis focuses on the deliveries in this mail center to evaluate the spatial coverage of drivers.

We use a local Nominatim[6] server to map GPS coordinates to street level location information. Our analysis show that the ratio of the streets with at least one delivery is 0.50, 0.44 and 0.33 in the three days covered in the dataset. The streets without a delivery are typically very short streets and/or with very small number residents. But, such streets are situated close to another street with a delivery.

Figure 2 (a) shows the CDF of the number of deliveries in a street in each day. In each three days, there are at most three deliveries in 90% of the streets. Similarly, Figure 2 (b) shows the distribution of the time difference between the first and last delivery on a street. In each of the three days, in about 85% of the streets, this duration is less than 1 hour. For a few streets, there are parcels delivered in the afternoon rounds that lead to large time span values.

These figures indicate that while the mobile workforce of the bpost provide a widespread spatial coverage, time spent in each street is usually only a fraction of the day as the time spent on a street does not span most of the day.

	Mean	Median	Standard Deviation
Round Duration	3.69 hrs	4.01 hrs	1.83 hrs
Number of Deliveries per Round	28.33	25	20.07
Time Between Two Deliveries	8.10 min	3.85 min	14.23 min
Total Distance	24.94 km	19.71 km	18.24 km
Distance Between Two Deliveries	770.85 m	421.1 m	1631.87 m

Table 1: Key Statistics on Postal Worker Behavior

We next turn our attention to some key statistics across all the delivery information in the entire dataset. The summary of these statistics is given in Table 1. Note that not all deliveries require a feedback of delivery information. Such parcels are not scanned and hence do not enter an event into the dataset. With this taken into account, these statistics are compatible with the values given to us by bpost.

2.3 Qualitative Analysis

In order to assess the workflow of the postal workers, we carried out an observational study with 6 postal workers in two urban settings, Brussels and Antwerp, and a rural area, Sint-Katelijne-Waver of Belgium. In each location, we accompanied two workers with a different mode of transportation.

We conducted a user shadowing study with *fly-on-the-wall* and *observation with inquiry* techniques [4]. In the first part of the tour,

[4] `http://corporate.bpost.be/about-us/bpost-at-a-glance?sc_lang=en`

[5] Note that this does not cover the entire bpost workforce but only those with the mobile terminal

[6] `http://wiki.openstreetmap.org/wiki/Nominatim`

we observed the postal workers without any comment or intervention. In the second part, we asked them questions on certain aspects related to their behavior. We used a head a mounted camera to capture a workers round in a video for later analysis. In addition, we used a GPS tracking application to track the route of the worker.

The mail rounds are broadly categorized into two groups based on the type of deliveries they make. In one group, workers deliver letters and small packets. In urban areas, workers in this group walk along their rounds with their caddies. In rural areas, they use small vehicles such as bikes and scooters to traverse extended distances but they still walk among a group of houses. We refer to workers in this group as *pedestrians*. In the second group, workers use motor vehicles to carry and deliver large parcels. In addition, they deliver letters for premium customers, i.e. people/firms that pay to receive their post before 9:00 am and high volume users, i.e. the entities that receive a large number of items in mail. These workers are referred to as *drivers*[7]. Drivers also distribute the letter bags of pedestrians to locations designated by the latter such as apartment buildings and depots, so they do not have to carry all the letters at all times. They pick up these bags left by the drivers along their rounds and deliver to their destinations. While a worker makes only one round in urban areas, in a rural area multiple rounds can be completed by the same worker using different modes of transportation, e.g. first on a car, then on a bike. In other words, a worker can be both a driver and a pedestrian during the course of the day. The summary of the mode of transportation of worker rounds is given in Table 2.

	Urban	Rural	Pedestrian (letters)	Driver (parcels)
Walk Only	✓		✓	
Car Only	✓			✓
Bike/Scooter + Walk		✓	✓	
Car + Walk (Two rounds)		✓	✓	✓

Table 2: Mode of Transportation for Postal Workers

Pedestrians perform their round in the morning, typically completing the delivery of letters by midday. Most of the parcels are also delivered in the morning by drivers. In addition, there are also afternoon driver rounds, delivering a small number of parcels. These drivers also collect outgoing letters in mailboxes placed around the city.

The round of a pedestrian is much shorter in comparison to that of a driver. They deliver letters to same people almost every day. Hence, a pedestrian establishes contacts with individuals on a regular basis and is able to build trust relationships with residents along that round. This gives them opportunities to collect voluntary information. On the other hand, a driver can only establish such contacts with the representatives of premium and high volume customers that they see every day. They also frequently interact with proxies who can provide collective information such as building concierges of large buildings since even though the addressee changes each time, a building with a large number of households is regularly to visited by the worker.

While a pedestrian does not have any time critical deliveries, drivers have a deadline for delivering mail for premium customers before 9:00 am. They also have a soft deadline for distributing pedestrian bags before 11:00 am. Still, both group of workers do not rush their deliveries. They take sufficient breaks along their round and have conversations with people in the vicinity. Drivers in cities are an exception to this as it may be difficult for them to find parking spots for their vehicles.

The workers have situational disadvantages for interacting with their mobile devices. Instead of single letters, pedestrians pick up a batch of letters from their caddies or bikes and distribute them to the mail boxes. Doing this, they leave their caddy or bike behind and hold the letters in their hands. Moreover, they cannot use their devices if/when they operate the bikes and scooters. Similarly, drivers cannot engage with mobile devices while they are behind the wheels and carry parcels from the vehicle to the destination address. Still, there are opportune moments to engage with the devices. For example, it is possible to check mobile devices after pedestrians complete a batch and while they return to their caddies. In drivers' case, most parcels require a signature and they have to wait for people to pick up the delivery after the bells are rung. They use their devices while waiting for the addressee to show up. All the workers noted that while it is possible they may not accept calls immediately, they get back to the caller as soon as they can.

	Pedestrian	Driver
Delivery with Deadlines		✓
Voluntary Opportunity	✓	
Proxy Interaction	✓	✓
Situational Disadvantages	✓	✓
Opportunity to Use Mobile Device	✓	✓

Table 3: Summary of Postal Worker Behavior

The summary of our observations is presented in Table 3.

2.4 Contextual Inquiry

After we had finished the qualitative analysis, we conducted a small semi-structured interview with the six postal workers we followed on their daily round. This group included four males (aged between 40 and 62) and two females (aged between 40 and 45). Two of these workers were merely on foot, two were only drivers, one of them performed a walking tour after finishing a driving round and one used a scooter on the round. The walk only and car only workers did rounds in urban settings whereas the other two had rounds in a rural area.

The goal of the interview is to find out if postal workers are indeed willing to take on the extra task of collecting information in their area and if so, what kind of tools and devices they would prefer to complete this task.

On the first question, whether they were willing to accept answering temporal-spatial questions when they are out on the street delivering letters and packages, all six persons responded that they would agree to do this. In addition, they said they could ask questions to residents they serve. More specifically, we asked whether they would feel comfortable asking people whether the area is perceived safe or dangerous. All subjects said they would not have any objections to such queries. Two works added that they do not see the value of asking this kind of questions, but they would do it anyway if requested by their management. None of the workers

[7]The previously studied dataset provides information about drivers

thought their primary task would prevent them from collecting such information.

We presented them a possible use case where they would receive a question at a certain geolocation during their tour. The question we suggested was the following: *Does the house at this address have solar panels on the roof?*. After explaining that this kind of question could cause harm to residents because failure to report a solar panel installation could result in paying a high penalty for tax avoidance, five of them responded that they would not endanger their social trust relation with the people they serve. One remarked that he would not want people to think he was a traitor.

Then, we inquired about the type of device they would like to use to collect this information, and we presented them with four possible choices: (a) smartphone, (b) smartwatch, (c) tablet, (d) in-house bpost scanning device. The most popular choice was the bpost terminal. The workers agreed in that they were not tech savvy. The drivers, in particular, noted that they would feel more comfortable with a device they already know and operate. Their second choice was to use a smartwatch, followed by the smartphone. None of the respondents opted the use of a tablet.

2.5 Summary

Our behavioral study shows that bpost workers have a widespread spatial coverage and a loose schedule. This suggests that these workers can be used to collect data about an urban area. Still, several factors need to be taken into consideration to design such a system. The factors can be summarized as follows.

- Workers have daily routines that they follow along their rounds. They can easily detect any changes in the area they serve.
- The workers enjoy a high degree of trust relationship with the public so they can get candid answers questions targeting residents. On the other hand, they are not willing to jeopardize this relationship.
- The postal workers frequently interact with citizens. While for some of them opportunities to collect information from volunteers, others can only acquire summarized collective information from proxies.
- While workers are able to engage with mobile devices along their rounds, they have situational disadvantages when they carry/hold items for interacting with devices. Their responses to the cues of the mobile device may be long depending on when they are notified.

3. SYSTEM

In this section, we introduce the basic structure of the system that oversees the collection of information. This system is built upon the our observations reported in the previous section. The system includes a smartwatch that extracts the context of the worker and decided when and where a question is presented to them.

Every task is associated with a query and an area of interest. Since the rounds of the workers are known in advance, it is possible to map a task to a particular worker. The assignment also takes worker behavior, interaction with other people and their responses (or non-responses) to the similar previous queries. This also removes the need for a live connection to the worker device. Instead, the worker synchronizes the device with the back in the beginning of the day to retrieve the assigned tasks and in the end of the day, uploads the answers to the queries to the servers.

The device checks whether it is inside one of the areas defined by the tasks. The sampling frequency for this check is low to save

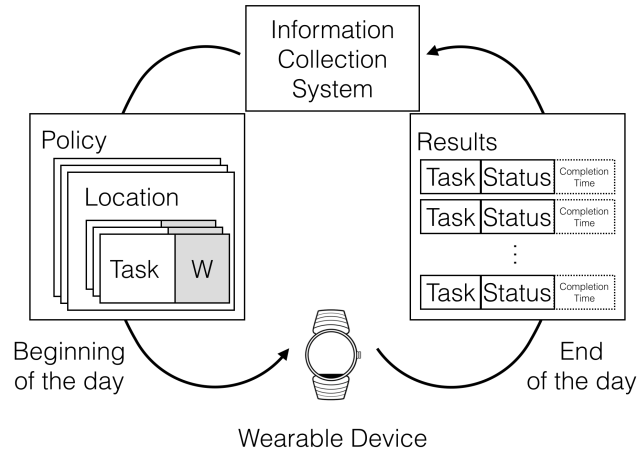

Figure 3: Synchronization of Task Policy and Results at the beginning and end of a working day between the device and the Information Collection System

battery power. Once it detects that it is inside one of the regions, the watch switches to the high-accuracy sensing mode that senses fine-granularity location information at a high rate. In this mode, it also activates various sensors it hosts such as an accelerometer, a microphone, etc to construct the context of the postal worker and to decide whether or not to present the query associated with the location.

Figure 3 shows the overview of the system architecture. The Information Collection System is a backend service that keeps a list of all tasks. Every task gets associated with a weight value that measures its importance for the service. In the beginning of the day, a worker synchronizes their device and downloads a *task policy* composed of a list of tasks organized by location. These tasks have previously been defined by one or more third parties known as *issuers*, and are characterized by a region, a query along with the weight. It also indicates whether the worker should answer question or direct it to another resident.

Once the worker starts its daily routine, the device tries to find the best moment to show a task request to the worker, based on its onboard sensors and processing capabilities. More in detail, the device can be in one of the following states:

Coasting Mode: This is the default mode of operation of a worker's device. In this mode, the device looks periodically at its current location using onboard geolocation sensors (e.g. GPS) at a low frequency (e.g. once per minute, or whatever low power mode is defined by the device's hardware) preserving the battery. This location is compared with the list of areas in the task policy. Whenever a device senses it is within one of the regions, it switches to *Sensing Mode*.

Sensing Mode: In this mode, the device analyzes worker context using a host of sensing capabilities. In particular, the *location sensor*, the *accelerometer* and *microphone* are used to detect whether the user is driving, walking, talking, or idle with any movement or speech[8].

By looking at the activity of the holder, the device tries to un-

[8]Modern mobile OSes provide APIs to retrieve such activities. It is also possible to compare the values retrieved from accelerometer with typical walking, driving speeds

Figure 4: User interface on the watch application

derstand what moment is best for presenting the question. The decision in this stage leverages our observations reported in the previous section. For a pedestrian the most appropriate moment is when the worker finished the current batch of letters and is on the way back to the caddy. For a driver, the device tries to understand whether the worker is waiting for an addressee to pick up the delivery. It also takes whether the question should be answered by the worker or a resident into account. Once such an opportune moment is detected, the device switches to *Task Mode*. If instead, the device detects to exit the area given in the task, it switches back to *Coasting Mode*.

Task Mode: In this mode, the worker is presented with the task that has the highest weight for the current location. Each task presents the user with a question on the mobile interface, and one or more possible answers (including free-form answers if specified in the policy). The worker can then provide an answer to the task marking it as *completed*, remove the task from the device interface and mark it as *rejected*, or simply not answering it, which marks the task as *ignored* after a timeout expires. Once a task is handled and marked, the device devices e whether to present the worker with another task for the same location, or to ignore all the other tasks for a given amount of time, then possibly switch back to Coasting Mode if the worker exited the current location range.

Once workers return to the mail center, they synchronize the devices for a second time, uploading the final status of the tasks to the Information Collection System. This includes the state of the task, i.e. whether it is completed, rejected or ignored, the answer selected by the worker if the task is completed and the time it took worker to complete or reject a task.

The system uses task state and completion time to evaluate the weight of similar tasks for a given worker in the following days.

Figure 4 shows the interface of the mobile smartwatch application presenting the query to a worker. The question is provided in a text-based interface that provides a multiple choice answer in a push/swipe menu.

4. DISCUSSION

We present a solution to the mobile crowdsourcing by leveraging a mobile workforce to address the lack of spatiotemporal coverage, accuracy and sustained engagement in traditional approaches to crowdsourcing. We use bpost postal workers as test subjects and study their daily workflow using quantitative and qualitative analyzing. Our analysis shows that using such a workforce is feasible for collecting qualitative information about country-wide points of interest.

In addition to shaping the information collection system, our observational study also sheds light on the types of crowdsourcing tasks a mobile workforce can be used for.

Data shows that parcel delivering postal workers who operate a motor vehicle deliver a parcel in approximately half of the streets in an urban area. In addition, workers who deliver letters on foot cover all the streets in the entire country and some of the streets are covered by drivers in the afternoon to pick up letters in the letter boxes. Hence, bpost workforce ensures a spacial coverage with a very high granularity for crowdsourcing applications. On the other hand, each street is visited at most three times a day. As a result, it is not possible to use such workforce for collecting real-time information. Instead, a postal workforce is valuable for collecting information that is not time-sensitive.

The crowdsourcing tasks that workers carry on can be both repetitive and one time. One time questions can be directly responded by the workers themselves. By leveraging their social status and their constant interaction with the residents they serve, workers can direct these questions to multiple people. Hence, the same question can be replied repetitively during the day.

While the workers are able to collect information from the residents by directing the queries to them, the nature of the questions in these queries should not endanger the relationship between the workers and the people. Hence, the tasks should not put the privacy of the people at risk. The trust relationship can be used to get candid responses from the public. However, this relationship should be treated delicately. Our interviews with the workers also highlight this notion as they are not willing to perform tasks that might be personal for residents.

We briefly describe the Information Collection System and the smartwatch application that is used for data collection. During the time this paper is written, this system was being deployed and tested. We plan to report its full-fledged evaluation in future work. The system is entirely automated in determining which query needs to be presented to the worker and when. The system may follow other approaches as well. The system can be fully manual where the worker decides when they can engage with the mobile device to introduce responses to queries. In this approach, the worker also decides what questions they answer. Between these two extremes, it is possible to follow a mixed solution where the worker manually indicates that they can handle a task. In such instances, the device then automatically presents the questions to the worker depending on their location, the time of the day and their context information that may include whether they are talking to someone else. Such selection may use the worker's history with the tasks that are similar to those in the device policy in resembling situations. These two approaches do not require the device to continuously track the location and hence introduces benefit in terms of battery. The fully manual approach on the other may cause them more time as they have to go through a list of questions to answer. Also, this approach may lead to a discrepancy where a question is answered many times whereas only a small number of responses are introduced to other queries. In future work, we will analyze these approaches more in detail and provide empirical comparison.

Using a mobile workforce such as postal workers can solve major problems experienced by traditional approaches to mobile crowdsourcing. By embedding such tasks to their daily work routines, we can achieve guaranteed spatial coverage without any dropouts and decreases in the engagement. It is not necessary to provide monetary incentives to information collectors. However, we plan to explore ways to incentivize the workers. For example, the workers may be more motivated if the providing responses to the queries become part of their key performance indicator and used for yearly bonuses and compensations. This can also contribute to the accuracy of the results and the sustained engagement with the crowdsourcing system.

This approach also faces a number of limitations. Our observations are specific to Belgian ecosystem. The worker-resident relationship may differ in other countries. In addition, the system is not suited to collect all kinds of tasks such as those that require real-time information flow and full temporal coverage.

This work has both practical and theoretical implications and calls attention to future urban designers, local authority and policymakers as well as social science researchers. From a practical standpoint, the suggested method to estimate the subjective well-being of a neighborhood can be transformed into a tool that can help urban designers to model the city in an informed manner. In terms of theory, our method can be used by social psychology and happiness economics practitioners to advance knowledge in their fields, and to expand their work across geographical boundaries.

Acknowledgement

IMEC has funded this research through ICON project SeRGIo.

5. REFERENCES

[1] F. Alt, A. S. Shirazi, A. Schmidt, U. Kramer, and Z. Nawaz. Location-based crowdsourcing: Extending crowdsourcing to the real world. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries*, NordiCHI '10, pages 13–22, 2010.

[2] D. Hristova, A. Mashhadi, G. Quattrone, and L. Capra. Mapping Community Engagement with Urban Crowd-Sourcing. In *Proc. When the City Meets the Citizen Workshop (WCMCW)*, June 2012.

[3] T. Kandappu, N. Jaiman, R. Tandriansyah, A. Misra, S.-F. Cheng, C. Chen, H. C. Lau, D. Chander, and K. Dasgupta. Tasker: Behavioral insights via campus-based experimental mobile crowd-sourcing. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '16, 2016.

[4] B. Martin, B. Hanington, and B. Hanington. *Universal Methods of Design: 100 Ways to Research Complex Problems, Develop Innovative Ideas, and Design Effective Solutions*. Rockport Publishers, 2012.

[5] J. Thebault-Spieker, L. G. Terveen, and B. Hecht. Avoiding the south side and the suburbs: The geography of mobile crowdsourcing markets. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '15, pages 265–275, 2015.