Nudge your Workforce: A Study on the Effectiveness of Task Notification Strategies in Enterprise Mobile Crowdsourcing

Sarah Bashirieh, Sepideh Mesbah, Judith Redi and Alessandro Bozzon
Delft University of Technology
Mekelweg 4
Delft, The Netherlands
sara.bashiri@gmail.com,s.mesbah@tudelft.nl
J.A.Redi@tudelft.nl,a.bozzon@tudelft.nl

Zoltán Szlávik and Robert-Jan Sips
IBM Benelux, Centre for Advanced Studies.
Johan Huizingalaan 765
Amsterdam, The Netherlands
zoltan.szlavik@ibm.nl.com
robert-jan.sips@ibm.nl.com

ABSTRACT

As crowdsourcing gains popularity, organisations seek ways to systematically and reliably involve their workforce with data processing pipelines. Mobile crowdsourcing allows for opportunistic task executions and thus, potentially, for higher throughput. However, how to engage and to retain employees in enterprise crowdsourcing campaigns is still an open research topic. This paper discusses the results of a study performed in IBM Benelux. We surveyed 93 employees to discover the factors that might affect engagement in mobile enterprise crowdsourcing. The survey informed the design of an experiment that aimed at investigating the effectiveness of different task notification strategies. We studied how factors such as time and context of notification can affect the participation and retention of employees. Results show that break times are the most suitable for crowd work, and that "aggressive" notification strategies act as deterrent for participation, while moderate yet regular nudges are the most likely to retain contributors.

CCS CONCEPTS

•**Information systems** →**Incentive schemes;** •**General and reference** →*Evaluation;* •**Applied computing** →*Enterprise information systems;*

KEYWORDS

Crowdsourcing; Notification; Enterprise; Workplace; Location-Aware Computing

1 INTRODUCTION

Enterprise crowdsourcing – i.e. the use of crowdsourcing approaches that harness the collective intelligence of an industrial firm's employees – is often seen as a way to systematically access the rich (and often tacit) knowledge and skills of a workforce across business divisions and hierarchical structures. Concrete examples of application domains for enterprise crowdsourcing include collaborative design and innovation, customers support, workplace awareness, and knowledge creation [7, 9, 16, 29, 34, 35].

The main advantage of enterprise crowdsourcing lies in the availability of a crowd of performers that are trusted, and operating in the context of an already retributed employment. Employees operate under contracts signed in order to commit to corporate norms and values, including, for instance, intellectual property rights [35]. The commitment from employees to employers helps decreasing malicious behaviours during task execution, and, arguably, makes employees always available for contribution, possibly in an opportunistic manner. To this end, mobile (or situated) crowdsourcing represent an interesting alternative w.r.t. traditional desktop crowdsourcing [19]. By exploiting the common availability of powerful mobile devices (e.g. smartphone), and considering that many people almost always have their devices with them, it is now possible to push tasks anywhere and anytime. Given the lack of agreed-upon regulations for the use of crowdsourcing within an organisation [5, 12], current crowd work models within enterprises must address delicate trade-offs, such as the one between an employee's primary work and the crowd work, or the one between the promotion of wide employee participation and management oversight and control over the workforce's activities [3].

With this work, we aimed at furthering the understanding on how enterprise mobile crowdsourcing (EMC) could be sustainably adopted in a traditional work environment. We investigated which factors might affect the engagement, participation, and retention of employees with EMC campaigns. The investigation took place in the multinational enterprise environment of *IBM Netherlands*, where we sought answer to the following research questions:

RQ1 : When and how would employees be willing to perform mobile crowdsourcing tasks in an enterprise environment?

The goal is to understand when employees are more likely to interact with their mobile devices, and which type of crowd work would they be willing to perform. Driven by previous literature, we created and advertised a survey that ultimately involved 93 participants.

We discovered that: 1) employees are generally willing to perform crowdsourcing tasks during some (but not all) of their break times. 2) Task duration is an important participation condition: results show that employees prefer to devote short attention spans to crowd work. 3) Personal motivations are preferential reasons for

Robert-Jan Sips is currently affiliated with *myTomorrows*. His current e-mail address is robert-jan.sips@mytomorrows.com.

UMAP'17, July 9–12, 2017, Bratislava, Slovakia

978-1-4503-4635-1/17/07...$15.00
DOI: http://dx.doi.org/10.1145/3079628.3079692

participation. For instance employees are willing to contribute if they sense the opportunity to learn, or to improve their workplace conditions.

To better understand how EMC could be sustainably adopted in an enterprise environment, we incorporated the results from the survey into the design of a mobile crowdsourcing experiment. The goal was to assess the effectiveness of alternative task notification strategies for enterprise crowdsourcing, thus, seeking answer to the following research question:

RQ2 How can notification (nudging) strategies affect the participation, retention and task commitment in enterprise mobile crowdsourcing?

We created a novel enterprise mobile crowdsourcing platform (*E-Crowd*) that allows employees to download and execute arbitrarily complex tasks directly onto their smartphone. The platform includes a mobile application with activity-awareness capabilities, to infer, for instance, whether a person is walking or not. Mechanisms devoted to the dynamic allocation of tasks, and push-based notifications, allowed us to experiment with time-aware and activity-aware notification strategies, while providing employees with a selection of different tasks to perform. Our findings suggest that task commitment (i.e. the amount of time devoted to EMC) is not affected by the selected notification strategy, while participation (i.e. the frequency of participation in EMC campaigns) and retention (i.e. the likelihood of an employee not to abandon the platform) are favoured by moderate yet regular nudges.

Enterprises are complex socio-technical systems. The role that enterprise crowdsourcing could play in traditional organisations is yet to be fully understood, let alone exploited. Our work contributes new insights on how and when employees can be engaged to contribute with crowd work. This additional knowledge can be used to promote the acceptance of this useful work paradigm within companies, and inform the design of a next generation of EMC systems that are efficient yet respectful of pre-existing work norms and dynamics.

2 RELATED WORK

Enterprise crowdsourcing. In enterprise crowdsourcing, organisations can benefit from the knowledge, intelligence and expertise of their employees to solve problems that cannot (or are not economical to) be completely outsourced to external workforces. Enterprise crowdsourcing can target workers external to the organisation [28] or a mix of external and internal populations [7]. The focus of this study is on enterprise crowdsourcing performed by a company's own workforce. Enterprise crowdsourcing proved applicable in a variety of domains, including collaborative design and innovation, customers support, workplace awareness, and knowledge creation [7, 10, 16, 29, 34, 35]. Yet, the theory and practice of Enterprise crowdsourcing currently lacks a clear understanding of how employees could be systematically engaged, and durably retained, in crowdsourcing campaigns. MCNet [26] addressed the crowdsourcing of WLAN measurements within an office environment, adducing as incentive for participation the increase of WLAN quality for employees. "Games for Crowds" [20] studies the effectiveness of a game-centred approach to Enterprise crowdsourcing, where users could play, create, and share simple games that harness the collective intelligence of employees within the enterprise. Stanculescu et al. [30] discuss the effectiveness of gamification techniques as engagement strategies for employees' participation in enterprise learning and social awareness tasks. MET [4] addressed the problem of providing the right incentives for Enterprise crowdsourcing participation, while preserving management oversight and control with monetary reward. To the best of our knowledge, our work is the first that studies how notification strategies can play a role in the success of enterprise crowdsourcing initiatives. We contribute an experiment based on a rigorous experimental design, and supported by a robust application-independent mobile crowdsourcing platform designed for the enterprise context.

Mobile Crowdsourcing in The Enterprise. Mobile phones provide flexibility in crowdsourcing task execution. Workers are able to ubiquitously contribute to crowdsourcing tasks using their mobile data connection, or local WiFi networks [2]. A key characteristic of mobile crowdsourcing relates to its opportunistic or participatory nature [15, 27]. Rather than sitting at a desk, workers can engage with a crowd work platform in a variety of different contexts and environments. This creates opportunities for more efficient execution of crowd work campaigns, but also provides a different set of challenges [21] (e.g. battery limitations, screen size, and lack of effective input devices). Several studies show that one of the main obstacles for large adoption is the presence of proper participation incentives [8, 27, 31, 36]. Commercial mobile crowdsourcing platforms provide monetary rewards, while academic and volunteer-based platforms rely on non-monetary incentives such as social rewards [17]. Few studies investigate the relationship between notification time, a user's location, and the likelihood and quality of users' responses. [23] show that user response probability to information notification is higher before and after a location change event. A recent work [22] investigates the issues of task notification design, to assess when, where, and to whom tasks should be suggested in on-the-go crowdsourcing. Results suggest that small changes in notification radius and timing can have a significant effect on individual participation and actions.

Novelty. In this paper, we combine qualitative and quantitative analysis to understand when and how employees would be willing to perform EMC tasks, and which notification strategies are the most effective in eliciting workers' contributions. We investigate a set of aspects relevant to enterprise mobile crowdsourcing (participation, retention and task commitment), that are only partially covered related work. No previous work investigating interruption and notification strategies [1, 6, 14, 25, 32], enterprise crowdsourcing [4, 30, 35], or mobile crowdsourcing [31] performed a survey, a longitudinal experiment, and post-experiment interviews in an enterprise environment. This is an important difference, as the demographics and motivation of participants are completely different, e.g. from those of a crowd of students.

3 INCENTIVES AND OPPORTUNITIES FOR EMPLOYEES' PARTICIPATION IN EMC CAMPAIGNS

The first part of our work investigates the likelihood of employees participating in enterprise mobile crowdsourcing (EMC) campaigns using their mobile phones (**RQ1**). At the same time, we seek insights into when employees are more likely to interact with their mobile phones – to check the presence of notifications from one or more installed applications. The assumption is that employees are normally busy with their business-related activities (e.g. meetings), and that unsolicited activities on their mobile devices can only happen opportunistically, or at pre-defined times of the day.

We created a survey[1], aimed at complementing results from previous work in terms of: 1) understanding of employees' motivations for (enterprise) crowd work; 2) knowledge of the amount of crowd work employees would be willing to execute; and 3) identification of the best times and context for effective notifications (nudging). The survey has been advertised using email invitations across departments, and the company's internal social network. An introductory page provided employees with information about the goals and scope of the survey, together with some explanation about crowdsourcing and its applications. The survey could be filled anonymously.

We targeted employees that used their smart phones as part of their daily activities. 93 employees participated in the survey; 20% of which were female, a ratio that reasonably approximates the company's gender distribution worldwide. 53% of the employees owned an iOS device, while 47% and Android device. The majority of participants (41%) were in the 46 to 55 year old range; 29% between 36 and 45; 19.2% between 26 and 35; 8.4% between 55 and 65; and 2.4% in the 18 to 25 year range. 42% of the participants worked in the engineering department, 41% in sales, and the rest from the other departments. 5% of the participants were managers.

3.1 How much crowd work would employees be willing to perform during a normal working day?

With this question, we aim at understanding how much crowd work (i.e. number of tasks) employees would be willing to execute each day during office hours. We rely on employees' self-assessment as it is not trivial to estimate an optimal number in an empirical manner, at least in a non-intrusive way [23]. As the duration of tasks can affect user voluntary participation, we also asked employees to express preferences about the preferred temporal length of execution. Table 1 reports the distribution of answers in the surveyed population. Numbers in bold highlight that the dominating factor is the **total** amount of daily time to be allocated (**from 5 to 10 minutes**), rather than the duration of single tasks, or the total amount of tasks. Employees indicated preference for short tasks, with length up to *2 minutes*. For instance 47% and 58% indicated the willingness to respectively perform work for up to 5 minutes (1 minute per task), or 10 minutes (2 minutes per task).

[1]The survey adopted in this work, as well as the complete set of results are available to readers at https://sites.google.com/site/enterprisemobilecrowd

	Duration		
# Daily Tasks	1 min	2 min	5 min
1 task	5.2%	17.1%	**54%**
2 to 5 tasks	**47%**	**58%**	**30.2%**
5 to 10 tasks	**26%**	17%	8%
10 to 15 tasks	12%	5.2%	0%
15 to 20 tasks	5.2%	0%	0%
More than 20	4%	2.6%	2.6%
No task	0%	0%	5.2%

Table 1: Preferred amount and duration of EMC tasks.

3.2 Which type of crowdsourcing tasks would incentivise employee participation?

Monetary incentives are often not suitable for enterprise crowdsourcing, both for managerial and legal reasons. We therefore investigate which non-monetary motivators explored in previous work (e.g. improving quality of work [26], improving the well-being of the workforce [29], or learning [30]) are the most popular among employees. *Learning Something New* was identified as the most popular motivation for participation (with 32% of preferences), followed by *Improving Work Conditions and Quality of Work* (24%) and *Having Fun* (18%). Interestingly, the *Improvement of the Company's performance* (14%) was perceived as more important than the improvement of one's *Performance Appraisal* (10%). No employee mentioned social relationships as the most motivating factor.

3.3 When do employees check the notifications on their smart phones?

Finally, we asked employees for information about their most likely availability during working time. Previous work [23] shows that employees are not responsive to the notifications on their phone at home and during meetings. By answering this question, we aim at pinpointing one or more time intervals, during a normal working day, when employees are more likely to check notifications.

We asked employees about whether they check notifications on their smartphones during coffee breaks, lunch breaks, during walking to the coffee machine, lunch room and meeting rooms. Results show that employees mostly check notifications on their smart phones *during standard breaks – excluding lunch time* – and *while walking* to coffee machines, lunch room and meeting rooms. This is an interesting outcome suggesting that: 1) break times have the potential to be good moments to solicit activities that are not related with employees' regular assignments; and 2) notification checking is more likely to happen while *walking* from one place to another (e.g. from the desk to the coffee machine) than when standing (e.g. at the coffee machine). The result also suggests that employees value their break times, possibly as a space for socialisation, and are therefore less willing to sacrifice it; therefore, the few minutes before and after the break are potentially more suitable for extemporary activity.

4 THE EFFECT OF NOTIFICATION STRATEGIES ON EMC PARTICIPATION, RETENTION AND TASK COMMITMENT

The second part of the work investigates how different notification strategies can affect the participation, retention and task commitment in enterprise mobile crowdsourcing (**RQ2**).

We designed and created an enterprise mobile crowdsourcing platform called *E-Crowd*, presented in Section 4.1. Inspired by the survey results reported in the previous section, we focused on two dimensions of study (independent variables), namely: 1) the *temporal distribution* of notifications; and 2) the *activity context* (i.e. walking) of employees. To reduce the space of experimental conditions, the type of available tasks were not considered as independent variables, but designed and implemented a fixed set of mobile crowdsourcing tasks.

Using a between-subject 2*2 factorial design, we created four distinct notification (nudging) strategies. The experimental design is discussed in the homonymous subsection. We made use of several metrics to assess (and compare) employees' *participation*, *retention* and *commitment* to the task; these metrics are described Section 4.2, while Section 4.3 presents the results of the experiment.

4.1 The E-Crowd platform

The platform consists of a back-end component and a front-end component. The back-end component (developed in *node.js*) includes functionality related to:1) registration and authentication of participants; 2) registration and allocation of crowdsourcing tasks; 3) integration and execution of custom notification strategies; 4) balanced allocation of participants to different experimental configurations; and 5) logging of participants activities. The front-end consists of an iOS application[2], implemented in *Swift 2*. Users can authenticate using their device Identifier (automatically retrieved from the phone metadata), so to allow anonymous participation and unique identification.

The application also includes activity-awareness capabilities, based on the *Core Motion* framework offered by the iOS environment. The motion activity is detected by the device hardware, specifically, the accelerometer and the gyroscope. The *Core Motion* framework provides an estimation of the current activity of the user, as inferred from the device.[3] To avoid battery draining (and, therefore, discourage employees' participation), motion activity is probed only at specific time spans (during break times) using the silent notification mechanism of iOS. [4]

Upon opening of the *E-Crowd* application, participants receive 5 tasks (Figure 1). To prevent bias in the results due to the order of task presentation, each participant is assigned a task list composed by a set of random tasks drawn from type-specific pools.

To minimise learning bias, executed tasks are never reassigned to users. By choosing a task, participants are provided with a brief description of the task's purpose, and with instructions for execution. Users can either "Start" o"Pass" a task (Figure 2). Starting a task allows users to execute the assigned activity and submit the results (Figure 3). When passing a task, participants are requested to provide a reason for rejection (e.g. *I don't like this task*, *The task is complicated*, *User Interface is not friendly* – see Figure 4). The logging functionality keeps track of several usage statistics, including task opening, task starting and passing, and task submission. The system also logs when the application is opened directly from the notification message. The application also checked at regular intervals if the participant decided to block notifications from the *E-Crowd* application.

4.2 Experimental design

The goal of the experiment was to measure how participation, retention, and task commitment (dependent variables) vary as a function of different notification configurations.

Independent Variables. The *Temporal Distribution* variable defines the time slots (during the working day) when the *E-Crowd* back-end pushes notifications to a given employee. The variable can assume two values, defined according to the survey results: 1) "Popular Break Time" (**PB**), which dispatches a single notification in each of the two time slots 10:00AM - 11.00AM (morning coffee break) and 12:00AM - 1.00PM (lunch break); and 2) "All Break Times" (**AB**), which includes also the time slots 9:00AM - 10.00AM (start of the day), 2:00PM - 3.00PM and 3:00PM - 4.00PM (afternoon coffee break) and 5:00PM - 6:00PM(end of the day).

The *Activity Context* variable could assume two values: 1) "Fixed Activities" (**FA)**, indicating that a notification will be sent to the participants exactly at the begin of break times (**AB** or **PB**); and 2) "Walking Activities" (**WA**), which send a notifications only after the detection of the *first* walking activity during the targeted time slot. Lack of motion detection, due, for instance, to the employee not moving from the desk (or not carrying the phone along) results in no emitted notification. Regardless the configuration, a welcome notification was delivered at the start of working days (8.50AM), to remind employees of the ongoing experiment.

	Group 1	Group 2	Group 3	Group 4
Popular **B**reaks	X		X	
All **B**reaks		X		X
Fixed **A**ctivities	X	X		
Walking **A**ctivities			X	X

Table 2: The four experimental treatment groups.

The resulting four treatment groups were organised as in Table 2. In *Group 1*, employees received notifications at popular break times; employees in *Group 2* received notifications at all break times, including the start and the end of a working day; in *Group 3*, employees received notifications at popular break times only if a walking activity is detected; finally, employees in *Group 4* received notification at all break times, but only when walking.

Our running hypothesis is that both the frequency of notification and the activity awareness would affect the likelihood of employees to: 1) react to notifications (thus opening the app); 2) execute crowd work; and 3) persist with their crowd work activities over time.

The between-subjects design was adopted to limit the effect of learning biases, but presents two main drawbacks: 1) it requires a large number of participants; and 2) it introduces variability due to the distribution of individual characteristics of the participants. The first issue has been tackled by deploying and advertising the tool with the company's internal network, which hosts a large potential number of employees, and resulted in a good participation. To minimise the chances of bias due to subjects' variability, at signup participants were randomly assigned to a treatment group.

[2]Due to technical constraints, this experiment includes only participants operating an iOS device.
[3]Examples of detected activities are walking, running, driving, or stationary. http://developer.apple.com/library/ios/documentation/EventHandling
[4]http://developer.apple.com/library/ios/documentation/NetworkingInternet

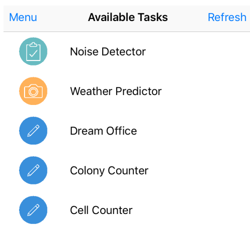

Figure 1: *E-Crowd*: Home Screen with Task List

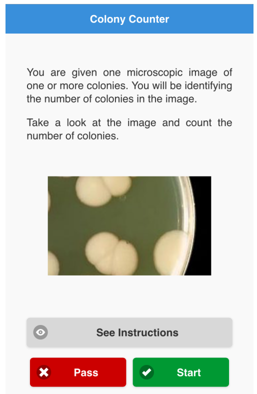

Figure 2: *E-Crowd*: Cell Annotation Task Instructions

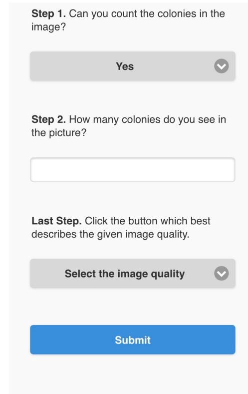

Figure 3: *E-Crowd*: Task Execution Interface

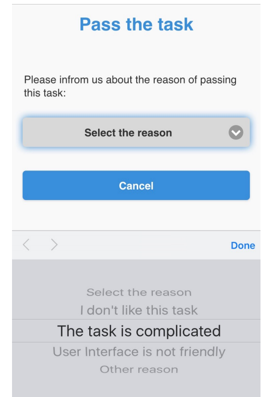

Figure 4: *E-Crowd*: Task Rejection Form

Dependent Variables. *Participation*, *retention* and *commitment* were quantitatively measured by analysing the logs produced by the *E-Crowd* application.

Participation is defined as the degree of involvement with the smartphone application; it is measured in terms of number of access to the application, and in terms of number of executed tasks. The average amount of application access and task execution falls into the category of *online behaviour metrics*, and are a good indicative of user participation, normalising for the sign up time.

To measure *retention*, i.e. the likelihood of an employee not to abandon the application, we compute at the proportion of participants who opened the mobile application more than once, and the frequency of their executed tasks. These numbers are captured in a *employee retention curve*, defined by the proportion of users who revisited *E-Crowd* and the frequency of their interactions [24]. The steepness of the slope indicates the level of engagement with the application. A steep slope implies that many employees interacted with the application only a few times; a flatter slope indicates that only a few employees abandoned the application after interacting with it just few times, while many employees had frequent interactions. Within each treatment group, we compute the number of access and executed task by all users and display a regression plot of the data.

Task commitment measures the amount of time dedicated to task execution; we consider it indication of employees' interest in the offered tasks, and a proxy for trustworthy execution. The time is measured as the amount of seconds between pressing the "Start" button, and successful task execution

Tasks Design. The survey highlighted a number of classes of tasks that are perceived by employees as motivating for participation. We included three types of tasks.

`T1 - Content Creation` (*Weather Predictor*) requires employees to take a picture of the outside weather in their current location. `T2 - Image Annotation` (*Colony Counter* and *Cell Counter*) requires employees to tap on their screen to *annotate* microscopic images of colonies or cells.

`T1` and `T2` tasks have direct benefit for the company (i.e. *improving the enterprise output*), as they are related to ongoing and advertised research endeavours.

`T3 - Survey` (*Dream Office* and *Noise Detector*) tasks are related to the *improvement of the quality of work*. They request employees to provide subjective assessments of environmental properties of the workplace, such as the experienced noise level, or the perceived office temperature and office lightning.

Tasks types were selected so to match both the preference of workers (from survey, *Dream Office* and *Noise Detector*) and the real world industrial needs of the company – specifically, the need to train the IBM Watson technology for application in the medical (*Colony Counter*, *Cell Counter*) and weather prediction domains. The selection of task types was also driven by the need to minimise the chances of confounding effects due to unbalance in task types, or due to individual task preferences.

All tasks were designed for optimal visualisation on a mobile phone screen. The amount work demanded by each task was tailored to fit a 1 minute execution time. To make ensure that employees were properly aware of the purpose of each task, the selection of a task from the list triggered a modal window containing a small explanation, always available in the task description page (see Figure 2). The time required to read the instructions of the task is considered a one-time penalty, and therefore negligible w.r.t. task execution time.

4.3 Experimental Results

The experiment has been conducted in the company's regional headquarter. Recruitment has been performed on a voluntary base, through advertisement. We sent emails in the corporate mailing list, advertised the experiment in the enterprise social network, and placed banners in strategic places (e.g. the cafeteria and bathrooms).

Participants interacted with the *E-Crowd* platform and mobile application for a number of days, depending on the signup date. The experiment ran for 2 months, between April 1st 2016 and June 1st 2016. During the observation interval, 83 employees installed the app. As employees could install the app at any time in this period, we considered the activities of each participant only during their first month of participation, regardless of their moment of enrolment. Notably, we logged no occurrence of employees disabling the notification functionality of the *E-Crowd* application.

4.3.1 General Participation Statistics. Participants differed per roles in the organisation, age groups and treatment groups. Figures are reported in Table 3–4. Most of the employees belonged to the *Engineering* and *Sales* department: this distribution reflects the actual allocation of employees within the company's headquarter. Other departments (e.g. legal) were also represented, but with fewer employees. We therefore aggregated them in a single category. Overall, managers accounted for 5% of the overall population. The distribution of employees across age groups also resembles the actual one of the company.

The amount of executed tasks is evenly distributed across departments. The differences across role groups and age groups are not significant (respectively $p > 0.19$ and $p > 0.28$). Participants in the `26--35` cohort are on average the most active. A consistent amount of participants (14%) were passive (i.e. inactive). 33% of them were managers. Due to technical limitations, we could not verify if these passive users uninstalled the application, or simply ignored all the incoming notifications.

We observed a considerable amount of activity during weekends, and outside official office hours (when participants did not receive notifications). As we are primarily interested in the direct effect of notification strategies, next sections will only report results based on activities and tasks completed during weekdays and office hours.

Dept.	$\#Em$	$\#Ex$	$\#Ex_{WH}$	$\%EX_{OH}$	μ_{ex}	$\mu_{ex_{WH}}$	$\%pas$
Engin.	24	304	239	21.3	2.6	9.96	12.5
Sales	32	370	191	48.3	11.5	5.97	15.6
Other	27	401	313	22	13.8	11.59	14.8

Table 3: Distribution of executed tasks across departments. $\#Em$: number of participants. $\#Ex$: number of executions including, off-time and weekends. $\#Ex_{WH}$: number of executions performed during working hours. $\%EX_{OH}$ is the percentage of task executions performed outside working hours. μ_{ex}: average amount of executions per users. $\mu_{ex_{WH}}$: average amount of executions per users during working hours. $\%pas$ indicates the percentage of passive users.

Age	$\#Em$	$\#Ex$	$\#Ex_{WH}$	$\%EX_{OH}$	μ_{ex}	$\mu_{ex_{WH}}$	$\%pas$
18-25	5	25	21	16	5	4.2	20
26-35	13	268	200	25.3	20.6	15.38	7.6
36-45	21	259	181	30	12.3	8.62	9.5
46-55	33	435	261	40	13	7.91	21.2
56-65	11	88	80	9	8	7.27	9

Table 4: Number of executed tasks in different age groups. Columns are the same as in Table 3.

Table 5 shows general participation statistics across the four treatment groups. Differences in the number of times the application has been opened are significant ($p < 0.004$), thus suggesting an effect due different treatment groups. **Group 1** (**PB**, **FA**) features the highest number of application opening *and* successful task executions, while **Group 3** (**PB**, **WA**) has the lowest.

	Group 1	Group 2	Group 3	Group 4
Participation Statistics				
# Open	319	102	85	149
# Pass	40	3	0	3
# Submit	259	188	114	136
Participation Statistics per Task Type				
T1 % pass	0.3%	0.5%	0%	0%
T1 % submit	6%	5.8%	6.1%	7.9%
T2 % pass	13%	0.5%	0%	2.2%
T2 % submit	53.2%	49.7%	59.6%	54.7%
T3 % pass	0%	0.5%	0%	0%
T3 % submit	27.5%	43%	34.5%	35.2%

Table 5: Participation statistics per treatment groups and task type. *#Open*: **Number of time that users open the app and receive a list of tasks.** *#Pass*: **Number of rejected tasks.** *#Submit*: **Number of executed tasks.**

We observe no statistically significant difference between groups in terms of amount of successfully submitted tasks ($p > 0.496$). Therefore, we cannot prove a relevant effect of the treatment groups on the number of successful executions. The distribution of submitted tasks also does not vary is a statistically significant manner across task types (see bottom of Table 5). Interestingly, the popularity of task types is aligned with previous findings in crowdsourcing literature (*Annotation* > *Survey* > *Content Creation*) [18]. To understand if the usability of the user interface could have been a motivating factor for inactivity or abandonment, we ran a parallel experiment where we asked participants to complete a Standard Usability Survey (SUS) [11]. The participants of this parallel experiment were drawn from the same population, but none was enrolled in the main experiment. The average SUS score for *E-Crowd* was found to be 72.5 (between "good" and "excellent"). The result shows that the application user interface did not suffer from usability problems that would significantly affect results.

4.3.2 How do temporal distribution and activity awareness of notifications affect the participation in EMC campaigns? We first investigate the effect that the amount of notifications might have on participation. Based on previous work [23], we hypothesise a negative effect on both the interaction with the app, and with the amount of executed tasks. To test the hypothesis, we compare these figures averaging across users assigned to the **PB**, **AB**, **FA**, and **WA** configurations. Table 6 and Table 7 show the results. Users assigned to a **PB** configuration opened the app, on average, 1.6 times more than **AB** users. A `Mann-Whitney test` ($U = 514$, $p = 0.21$) shows that it is not possible to account this difference to varying temporal distributions. Users without activity awareness (**FA**) opened the application more often, but the average difference is not statistically significant (`Mann-Whitney` $U = 471.5$, $p = 0.11$). Differences in terms of number of executed tasks are minimal (Table 7) and not statistically significant across temporal distributions (**PB** vs. **AB**: `Mann-Whitney`, $U = 583.5$, $p = 0.16$). Such differences are instead greater (and statistically significant) across activity awareness configurations (**FA** vs. **WA**: `Mann-Whitney U-test`, $U = 662.5$, $p = 0.02$).

	$\#Open$	$\#Em$	μ_{op}	σ_{op}	Med_{op}
PB	397	35	11.34	26.09	4
AB	243	33	7.36	14.77	4
FA	419	38	11.02	25.12	4.5
WA	221	30	7.96	15.32	4

Table 6: App opening statistics. ***#Open*: number of app opening.** ***#Em*: number of participants.** **μ_{op}, σ_{op}, Med_{op}: average, std.deviation, and median # of app opening per user.**

	$\#Submit$	$\#Em$	μ_{ex}	σ_{ex}	Med_{ex}
PB	373	33	11.30	13.69	8
AB	324	31	10.45	12.66	6
FA	447	34	13.14	15.40	8
WA	250	30	8.33	9.50	5.5

Table 7: Task execution statistics . ***#Submit*: number of submitted tasks.** ***#Em*: number of participants.** **μ_{op}, σ_{op}, Med_{op}: average, std.dev., and median # of app opening per user.**

While we can't provide strong statistical evidence for our initial hypothesis, the result suggests that higher frequency of notifications (during planned break times) might be beneficial in terms of interaction with the applications, but not in terms of the amount of executed crowd work, which tends to be reduced. On the other hand, a reduction on the number of notifications due to activity awareness has negative effects on both interaction with the application and with the number of executed tasks.

4.3.3 How do temporal distribution and activity awareness of notifications affect user retention in EMC campaigns? We measure user retention curve by calculating the amount of users who executed exactly $\boldsymbol{x}$ (frequency of their executed tasks) tasks in total [24]. The steepness of the slope is an indicator of the level of user retention. In the light of the results obtained with user participation, we hypothesise that an excessive amount of notification might have a detrimental effect on users, which tend to abandon the application after just executing a few tasks.

Figure 5 shows the user retention curves across experimental configurations. In terms of number of interactions with the application, the rate of change of the **AB** configuration is double with respect to the **PB** configuration (**PB** $coeff = -0.079$, $intercept = 3.281$, **AB** $coeff = -0.18$, $intercept = 4.606$). The difference is even more pronounced when considering activity awareness (**FA** $coeff = -0.092$, $intercept = 3.282$, **WA** $coeff = -0.233$, $intercept = 4.75$). In terms of number of submitted tasks, the rate of change of the **AB** configuration is approximately 20% larger than the one of the **PB** configuration (**PB** $coeff = -0.027$, $intercept = 2.48$, **AB** $coeff = -0.35$, $intercept = 2.47$). The rate of the **WA** is, in absolute terms, the larger (**FA** $coeff = -0.022$, $intercept = 2.00$, **AB** $coeff = -0.41$, $intercept = 2.62$). The **AB** and **WA** configurations feature steeper slopes.

These trends are consistent with the results from the user participation analysis. The results supports our previous conclusion: a wrong amount of notifications (respectively, too many in the **AB** configuration, or too few in the **WA** configuration) can lead to a more "enthusiastic" initial participation, but could also cause a quicker *crowd out* effect.

4.3.4 How do temporal distribution and activity awareness of notifications affect task commitment in EMC campaigns? Task commitment is measured in terms of the amount of time dedicated to complete the execution of a task. Considering the previous results, we hypothesis commitment to also be affected by variations in notifications configuration. Table 8 shows the average task duration time. Results contradict our hypothesis, as that task commitment is not affected by the temporal distribution of notifications (**PB** vs. **AB**: `Mann-Whitney U-test`, $U = 492$, $p = 0.399$) nor by the presence of activity awareness functionalities. (**FA** vs. **WA**: `Mann-Whitney U-test`, $U = 437$, $p = 0.164$).

	μ	σ	$Median$
PB	50.45	24.05	40.33
AB	47.80	27.64	43.6
FA	48.44	29.62	39.7
WA	50.00	20.8	45.66

Table 8: Mean, standard deviation and median of response time (seconds) to complete tasks.

5 DISCUSSION

The survey (**RQ1**) provided a number of relevant insights. Results show that employees are not bothered by smartphone notifications, at least during popular breaks. Employees already have non-work related interactions with their smartphone during coffee breaks, while walking to lunch, walking to coffee breaks and while walking to a meeting room. Note that all these activities include walking. These results suggest that all these moments could be potentially suitable for crowd work execution, especially if the tasks to be executed are of short duration: employees declared a preference for short tasks, with the duration of up to two minutes. The nature of the task is also of importance: *learning*, *purpose*, and *well-being* appear to be the most compelling reason for participation.

Answering **RQ1** informed the design of an experiment that accounted for the preferences of the targeted population. We studied notification strategies on mobile phones as a mean to pro-actively invite employees at strategic times (**RQ2**). The strategies addressed in this work had no significant effect in terms of commitment (i.e. average amount of time spent executing each task), thus providing additional evidences of the general positive attitude of employees towards the crowd work, once engaged with it. Strategies were instead influential for participation and retention: frequent notifications lead to lower retention and task executions. The result suggests that, in EMC, an excessive amount of notification can facility the feared "crowd-out" effect.

The presence of activity-awareness functionalities resulted in significantly lower participation and retention. This is an interesting result, that contrasts with the outcome of the survey. It is possible to hypothesise that the presence of notifications on the mobile device (e.g. after returning from a break, if the device has been left on the desk) is more important than their real-time detection. The hypothesis is consistent with recent findings [22] that show how (in the context of situational crowdsourcing) small changes in notification timing can have a significant effect on individual participation and actions. The verification of this hypothesis in the context of EMC is left to future work.

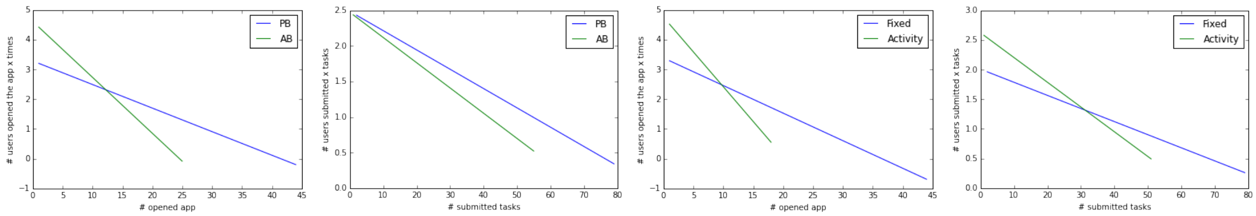

(a) #App Opening (PB vs AB) **(b) #Submitted Tasks. (PB vs AB)** **(c) #App Opening. (FA vs WA)** **(d) #Submitted Tasks. (FA vs WA)**

Figure 5: User retention curves for number of app openings and number of submitted tasks.

Regardless of treatment groups and strategies, a larger number of tasks were executed between `9:00AM to 10:59AM` comparing to other time slots. This shows that users are more willing to execute tasks in the morning. This result is in line with the survey results. The experimental observation interval had a duration of 4 weeks, i.e. 22 working days. Considering all active users, the average amount of execution per user is 10.4 during 22 days, meaning 0.47 task per day. This average result can be considered satisfactory, given the adopted advertisement and recruitment strategy. However, it is not in line with the outcome of the survey. Due to the anonymity of participants, we cannot exclude that external factors (e.g. pressing deadlines) could have affected participation.

Despite notifications being blocked outside office hours, 30% of all executions were performed during time slots that could be considered as belonging to employees' "free" time. This is an interesting result, that we further investigated with informal interviews conducted (at the end of the experimental period) with employees that decided to disclose their identity. These participants revealed that they were indeed interested in contributing to the objectives of the tasks offered by *E-Crowd*, but that they only had time to use the app during weekends, and outside office hours. These interviews provide evidence that, when properly motivated, EMC can connect employees with relevant tasks also during their free time.

Threats to validity. To minimise internal validity threats, we considered issues such as *history*, *selection*, *instrumentation* and *maturation* [13]. The *history* effect is concerned with the possibility that participants discuss the tasks among themselves, or find out about their special treatment group. Such an effect cannot be ruled out, but its impact is minimised as all treatment groups were similarly affected. *Selection* threat was addressed by randomly assigning participants to treatment groups. Moreover, the demographic distribution of participants in both the survey and the experiment resembled the one of the company. The limitation to participants operating an iOS device could have introduced a selection bias that we cannot rule out. The *instrumentation* threat was addressed by using the same procedure to measure the dependent variables for all the treatment groups. By relying on activity awareness functionalities that are built-in iOS devices, we minimised the likelihood of measurement errors due to implementation issues. To account for the *maturation* threat, we randomly created task lists for each participants, and considered only the first month of the contributions for all treatment groups.

Based on the above discussion, we believe the results provided in this work to be valid contributions.

6 CONCLUSION

Our research aims at better understanding the opportunities and limitations with the application of a mobile crowdsourcing paradigm in the enterprise environment. In this paper, we have presented the results of a survey and an experiment conducted in *IBM Netherlands*. Being the first experiment of this kind performed in an enterprise environment, the results are intrinsically important and novel.

The survey provided novel insights about: 1) the time slots (e.g. breaks) and context (e.g. while walking) when employees are more likely to interact with their mobile devices and perform crowd work; and 2) the type of crowd work that employees would be willing to perform during a working day. The results of our survey confirm the result of other exploratory studies performed in other environments and companies, but give additional insights on the types and duration of tasks preferred by employees.

A novel enterprise mobile crowdsourcing platform (`E-Crowd`) enabled the implementation of notification strategies with varying temporal distribution and activity awareness. The platform allowed for the execution of an experiment that involved 83 employees. We found that timely notifications can foster participation and retention. We found significant differences among treatment groups in terms of participation and retention. The outcome of the experiment is in contrast with the survey results, as activity awareness lead to lower participation and retention.

Future Work. This work provides plenty of inspirations for future research directions. We plan to further investigate the potential benefit that *learning* tasks could provide in terms of user engagement and retainment. Applications such as DuoLingo [33] clearly demonstrated that significant crowdsourcing results can be obtained when the incentives of the user and the one of the company align.

As 30% of all executions were voluntarily performed outside office hours, we are interested in understanding the potential for EMC when employees are not at work. The research will necessarily address ethical concerns about whether organisations should be allowed to enable EMC outside working hours, thus invading employees' private time. It will be interesting to assess the benefits, in terms of amount and quality of work, of EMC tasks performed outside office hours. Other relevant directions include the development and testing of personalised task notification strategies, and the potential benefits deriving from the adoption of incentive schemes based on gamification techniques.

REFERENCES

[1] Piotr D. Adamczyk and Brian P. Bailey. 2004. If Not Now, when?: The Effects of Interruption at Different Moments Within Task Execution. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 271–278. DOI: http://dx.doi.org/10.1145/985692.985727

[2] Florian Alt, Alireza Sahami Shirazi, Albrecht Schmidt, Urs Kramer, and Zahid Nawaz. 2010. Location-based crowdsourcing: extending crowdsourcing to the real world. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries*. ACM, 13–22.

[3] Anupriya Ankolekar, Filippo Balestrieri, and Sitaram Asur. 2016. MET: An Enterprise Market for Tasks. In *Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion (CSCW '16 Companion)*. ACM, New York, NY, USA, 225–228. DOI: http://dx.doi.org/10.1145/2818052.2869113

[4] Anupriya Ankolekar, Filippo Balestrieri, and Sitaram Asur. 2016. MET: An Enterprise Market for Tasks. In *Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion*. ACM, 225–228.

[5] Obinna Anya. 2015. Bridge the Gap!: What Can Work Design in Crowdwork Learn from Work Design Theories?. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. ACM, New York, NY, USA, 612–627. DOI: http://dx.doi.org/10.1145/2675133.2675227

[6] Brian P Bailey and Joseph A Konstan. 2006. On the need for attention-aware systems: Measuring effects of interruption on task performance, error rate, and affective state. *Computers in human behavior* 22, 4 (2006), 685–708.

[7] Osvald M Bjelland and Robert Chapman Wood. 2008. An inside view of IBM's' Innovation Jam'. *MIT Sloan management review* 50, 1 (2008), 32.

[8] Alessandro Bozzon, Marco Brambilla, Stefano Ceri, Andrea Mauri, and Riccardo Volonterio. 2014. *Pattern-Based Specification of Crowdsourcing Applications*. Springer International Publishing, Cham, 218–235. DOI: http://dx.doi.org/10.1007/978-3-319-08245-5_13

[9] Alessandro Bozzon, Ilio Catallo, Eleonora Ciceri, Piero Fraternali, Davide Martinenghi, and Marco Tagliasacchi. 2012. A Framework for Crowdsourced Multimedia Processing and Querying. In *CrowdSearch 2012 workshop at WWW 2012*, ceur ws (Ed.), Vol. 842. 42–47. http://ceur-ws.org/Vol-842/crowdsearch-bozzon.pdf

[10] Alessandro Bozzon, Piero Fraternali, Luca Galli, and Roula Karam. 2014. Modeling crowdsourcing scenarios in socially-enabled human computation applications. *Journal on Data Semantics* 3, 3 (2014), 169–188. http://dx.doi.org/10.1007/s13740-013-0032-2

[11] John Brooke and others. 1996. SUS-A quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7.

[12] Melissa Cefkin, Obinna Anya, Steve Dill, Robert Moore, Susan Stucky, and Osariemo Omokaro. 2014. Back to the Future of Organizational Work: Crowdsourcing and Digital Work Marketplaces. In *Proceedings of the Companion Publication of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW Companion '14)*. ACM, New York, NY, USA, 313–316. DOI: http://dx.doi.org/10.1145/2556420.2558858

[13] Thomas D Cook, Donald Thomas Campbell, and Arles Day. 1979. *Quasi-experimentation: Design & analysis issues for field settings*. Vol. 351. Houghton Mifflin Boston.

[14] Mary Czerwinski, Ed Cutrell, and Eric Horvitz. 2000. Instant Messaging: Effects of Relevance and Timing. In *People and Computers XIV: Proceedings of HCI 2000*, Vol. 2. 71–76. https://www.microsoft.com/en-us/research/publication/instant-messaging-effects-of-relevance-and-timing/

[15] Vincenzo Della Mea, Eddy Maddalena, and Stefano Mizzaro. 2015. Mobile crowdsourcing: four experiments on platforms and tasks. *Distributed and Parallel Databases* 33, 1 (2015), 123–141. DOI: http://dx.doi.org/10.1007/s10619-014-7162-x

[16] Stephen Dill, Robert Kern, Erika Flint, and Melissa Cefkin. 2013. The Work Exchange: Peer-to-Peer Enterprise Crowdsourcing.. In *HCOMP (Works in Progress / Demos) (AAAI Workshops)*, Vol. WS-13-18. AAAI. http://dblp.uni-trier.de/db/conf/hcomp/hcomp2013w.html#DillKFC13a

[17] Johann Füller, Katja Hutter, and Mirijam Fries. 2012. Crowdsourcing for Goodness Sake: Impact of Incentive Preference on Contribution Behavior for Social Innovation. *Adv Int Market* 23 (2012), 137–159.

[18] Ujwal Gadiraju, Ricardo Kawase, and Stefan Dietze. 2014. A Taxonomy of Microtasks on the Web. In *Proceedings of the 25th ACM Conference on Hypertext and Social Media (HT '14)*. ACM, New York, NY, USA, 218–223. DOI: http://dx.doi.org/10.1145/2631775.2631819

[19] Jorge Goncalves, Simo Hosio, Vassilis Kostakos, Maja Vukovic, and Shin'ichi Konomi. 2015. Workshop on Mobile and Situated Crowdsourcing. In *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers (UbiComp/ISWC'15 Adjunct)*. ACM, New York, NY, USA, 1339–1342. DOI: http://dx.doi.org/10.1145/2800835.2800966

[20] Ido Guy, Anat Hashavit, and Yaniv Corem. 2015. Games for crowds: A crowdsourcing game platform for the enterprise. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. ACM, 1860–1871.

[21] Jiyin He, Kai Kunze, Christoph Lofi, SanjayK. Madria, and Stephan Sigg. 2014. Towards Mobile Sensor-Aware Crowdsourcing: Architecture, Opportunities and Challenges. In *Database Systems for Advanced Applications*, Wook-Shin Han, Mong Li Lee, Agus Muliantara, Ngurah Agus Sanjaya, Bernhard Thalheim, and Shuigeng Zhou (Eds.). Lecture Notes in Computer Science, Vol. 8505. Springer Berlin Heidelberg, 403–412. DOI: http://dx.doi.org/10.1007/978-3-662-43984-5_31

[22] Yongsung Kim, Emily Harburg, Shana Azria, Aaron Shaw, Elizabeth Gerber, Darren Gergle, and Haoqi Zhang. 2016. Studying the Effects of Task Notification Policies on Participation and Outcomes in On-the-go Crowdsourcing. In *Proceedings of the Fourth AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2016*.

[23] Vlaho Kostov, Takashi Tajima, Eiichi Naito, and Jun Ozawa. 2006. Analysis of appropriate timing for information notification based on indoor user's location transition. In *Pervasive Computing and Communications, 2006. PerCom 2006. Fourth Annual IEEE International Conference on*. IEEE, 6–pp.

[24] Edith Law and Luis Von Ahn. 2009. Input-agreement: a new mechanism for collecting data using human computation games. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1197–1206.

[25] Tadashi Okoshi, Julian Ramos, Hiroki Nozaki, Jin Nakazawa, Anind K. Dey, and Hideyuki Tokuda. 2015. Reducing Users' Perceived Mental Effort Due to Interruptive Notifications in Multi-device Mobile Environments. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 475–486. DOI: http://dx.doi.org/10.1145/2750858.2807517

[26] S. Rosen, Sung-Ju Lee, Jeongkeun Lee, P. Congdon, Z.M. Mao, and K. Burden. 2014. MCNet: Crowdsourcing wireless performance measurements through the eyes of mobile devices. *Communications Magazine, IEEE* 52, 10 (October 2014), 86–91. DOI: http://dx.doi.org/10.1109/MCOM.2014.6917407

[27] John P Rula, Vishnu Navda, Fabián E Bustamante, Ranjita Bhagwan, and Saikat Guha. 2014. No one-size fits all: Towards a principled approach for incentives in mobile crowdsourcing. In *Proceedings of the 15th Workshop on Mobile Computing Systems and Applications*. ACM, 3.

[28] Ada Scupola and Hanne Westh Nicolajsen. 2014. The impact of enterprise crowdsourcing on company innovation culture: The case of an engineering consultancy. In *Nordic Contributions in IS Research*. Springer, 105–120.

[29] Robert-Jan Sips, Alessandro Bozzon, Gerard Smit, and Geert-Jan Houben. 2015. *Engineering the Web in the Big Data Era: 15th International Conference, ICWE 2015, Rotterdam, The Netherlands, June 23-26, 2015, Proceedings*. Springer International Publishing, Cham, Chapter The Inclusive Enterprise: Vision and Roadmap, 621–624. DOI: http://dx.doi.org/10.1007/978-3-319-19890-3_43

[30] Laurentiu Catalin Stanculescu, Alessandro Bozzon, Robert-Jan Sips, and Geert-Jan Houben. 2016. Work and Play: An Experiment in Enterprise Gamification. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing, CSCW 2016, San Francisco, CA, USA, February 27 - March 2, 2016*. 345–357. DOI: http://dx.doi.org/10.1145/2818048.2820061

[31] Rannie Teodoro, Pinar Ozturk, Mor Naaman, Winter Mason, and Janne Lindqvist. 2014. The Motivations and Experiences of the On-demand Mobile Workforce. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. ACM, New York, NY, USA, 236–247. DOI: http://dx.doi.org/10.1145/2531602.2531680

[32] Liam D. Turner, Stuart M. Allen, and Roger M. Whitaker. 2015. Interruptibility Prediction for Ubiquitous Systems: Conventions and New Directions from a Growing Field. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 801–812. DOI: http://dx.doi.org/10.1145/2750858.2807514

[33] Luis von Ahn. 2013. Duolingo: Learn a Language for Free While Helping to Translate the Web. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces (IUI '13)*. ACM, New York, NY, USA, 1–2. DOI: http://dx.doi.org/10.1145/2449396.2449398

[34] M. Vukovic. 2009. Crowdsourcing for Enterprises. In *2009 Congress on Services - I*. 686–692. DOI: http://dx.doi.org/10.1109/SERVICES-I.2009.56

[35] Maja Vukovic and Claudio Bartolini. 2010. Towards a research agenda for enterprise crowdsourcing. In *Proceedings of the 4th International Conference on Leveraging Applications of Formal Methods, Verification, and Validation - Volume Part I (ISoLA'10)*. Springer, Berlin, Heidelberg, 425–434. http://dl.acm.org/citation.cfm?id=1939281.1939323

[36] Jie Yang, Alessandro Bozzon, and Geert-Jan Houben. 2015. Harnessing Engagement for Knowledge Creation Acceleration in Collaborative Q&A Systems. In *User Modeling, Adaptation and Personalization: 23rd International Conference, UMAP 2015, Dublin, Ireland, June 29 – July 3, 2015. Proceedings*, Francesco Ricci, Kalina Bontcheva, Owen Conlan, and Séamus Lawless (Eds.). Springer International Publishing, Cham, 315–327. DOI: http://dx.doi.org/10.1007/978-3-319-20267-9_26