

Exploring Crowdsourced Work in Low-Resource Settings

Manu Chopra
Microsoft Research India

Indrani Medhi Thies
Microsoft Research India

Joyojeet Pal
Microsoft Research India

Colin Scott*
Microsoft Research India

William Thies
Microsoft Research India

Vivek Seshadri
Microsoft Research India

ABSTRACT

While researchers have studied the benefits and hazards of crowdsourcing for diverse classes of workers, most work has focused on those having high familiarity with both computers and English. We explore whether paid crowdsourcing can be inclusive of individuals in rural India, who are relatively new to digital devices and literate mainly in local languages. We built an Android application to measure the accuracy with which participants can digitize handwritten Marathi/Hindi words. The tasks were based on the real-world need for digitizing handwritten Devanagari script documents. Results from a two-week, mixed-methods study show that participants achieved 96.7% accuracy in digitizing handwritten words on low-end smartphones. A crowdsourcing platform that employs these users performs comparably to a professional transcription firm. Participants showed overwhelming enthusiasm for completing tasks, so much so that we recommend imposing limits to prevent overuse of the application. We discuss the implications of these results for crowdsourcing in low-resource areas.

CCS CONCEPTS

• **Social and professional topics** → **Employment issues**;
• **Human-centered computing** → *Text input*; *Smartphones*.

KEYWORDS

Mobile Crowdsourcing; Digitization; Regional Languages; Microtasks; Low-resource Settings

ACM Reference Format:

Manu Chopra, Indrani Medhi Thies, Joyojeet Pal, Colin Scott, William Thies, and Vivek Seshadri. 2019. Exploring Crowdsourced Work

*Colin Scott is currently part of the Next Billion Users group at Google.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI 2019, May 4–9, 2019, Glasgow, Scotland Uk

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300611>

in Low-Resource Settings. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland Uk*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3290605.3300611>

1 INTRODUCTION

Paid crowdsourcing platforms such as Amazon Mechanical Turk [1] have seen high representation of workers from India [46]. To date, most of these workers have been relatively well educated, fluent in English, urban, and have access to high speed internet [36]. However, we believe that workers from rural India also have much to gain, and much to contribute, as participants in crowdsourcing platforms. Given that 75% of rural Indians live on less than INR 33 (USD 0.5) a day [10], paid digital work could be a good source of supplementary income. It could also bolster digital literacy and skills, potentially unlocking other earning opportunities as well as broader interest or participation in positive online communities. Finally, crowdsourced work in local languages has the potential to bring marginalized cultural artifacts into the digital realm.

Despite these potential benefits, people in rural India are largely excluded from crowdsourced digital work. First, most employers do not choose to send work to this population unless they can receive the quality and pricing that is competitive with market alternatives. In fact, companies and research organizations that need crowd work have largely overlooked rural citizens, arguably, due to market factors such as lack of trust in rural labor. Second, it has not been clear that there is a scalable source of valuable tasks that could be completed with only local-language skills. Finally, as current crowdsourcing platforms generally require a computer and Internet connection, it is not clear whether constraints of infrastructure and technical expertise would prevent platforms from reaching rural areas.

Recently, however, we believe that there is a new window of opportunity to address the challenges mentioned above. In addition to the growing penetration of smart phones and Internet connectivity, a key part of our thinking is to envision the government as a new source of crowdsourced tasks. For example, the Digital India mission [39] has mandated digitization of all government documents. Such documents are often handwritten in one of India's 120+ local languages, making them unsuitable for off-the-shelf OCR technology

and a good match for the skills of local populations. The Indian government has also demonstrated unusual willingness to pay for work in rural areas. For example, the Mahatma Gandhi National Rural Employment Guarantee Act (MGN-REGA or NREGA) guarantees work to any rural household that requests it [43]. The scheme has paid over 6.4 billion dollars to 77 million people in 2017 [43]. Until now, programs like NREGA have focused on manual unskilled labor in support of local infrastructure (such as roads). In the future, could such a welfare program be broadened to include digital tasks such as document digitization, creating value not only for the government but also building digital literacy skills, and potential upward mobility, among rural participants?

As a first step towards answering this question, this paper explores whether it is possible for workers in rural India to digitize handwritten local-language text on a smartphone. Our goal is both to study the technical feasibility of such work—i.e., the accuracy and cost of text input as compared to prevalent transcription services—as well as the qualitative experiences of the workers.

To this end, we built a prototype Android application that enables digitization of handwritten text. We focus on text written in Devanagari script [17], used by many Indian languages, which workers input using the Swarachakra keyboard [8, 35]. We conducted an initial two-week user study with 12 participants in Amale, a village in the state of Maharashtra, India. This study was used to refine our application and was followed by a controlled two-week study with 32 participants in Soda, in the state of Rajasthan. Both these villages are heavily resource-constrained [9, 27]. Each participant in Amale digitized up to 8000 pseudo-handwritten Marathi words and each participant in Soda digitized up to 6000 real handwritten Hindi words. In order to explore the market viability of the platform, we offered compensation that was commensurate with local wages for semi-skilled labor. This translated to an honorarium of INR 3000 (about USD 45) per participant. At the end of the study, we interviewed all our participants to understand their experience.

Our study reveals the following findings. First, our study shows that, in addition to learning to use the Swarachakra keyboard, users showed both the ability and willingness to complete digitization tasks over two weeks. Second, users achieved accuracy high enough to perform comparably to a professional transcription firm in a crowdsourced setting. Our results show that a crowdsourcing platform employing our users can achieve an accuracy of 98.9% at a cost of INR 1.19 per word. In comparison, a market alternative achieved an accuracy of 98.4% at a cost of INR 1.00 per word on the same dataset. Third, every user stated that they were not aware of the concept of digital work, that they found the work to be rewarding and fun, and that they were fairly

compensated for their work. In some cases, users' enthusiasm for the work also led to unusual working hours, which underlines the already-recognized need for crowdsourcing platforms to protect the best interests of participants, both from requesters and from themselves. Finally, users reported that they got better at typing in Hindi, so much that they started using Hindi in messaging apps like WhatsApp.

Our paper takes a measurable step in making crowdsourcing platforms more inclusive of workers from rural India. To this end, our paper focuses on text input, and is the first work to measure the accuracy with which non-English speaking rural workers with limited exposure to smartphones can digitize handwritten text in regional languages. Our results show that rural workers can achieve sufficiently high quality of work to participate in a paid crowdsourcing platform. There are additional challenges associated with reaching and training rural workers in making such a platform viable. With the proliferation of smartphones and growing access to cellular data, we believe future work can build scalable solutions to address these challenges.

2 RELATED WORK

2.1 Future of Work

Our work is situated within a vibrant global conversation around the future of work, and how technologies such as crowdsourcing could offer benefits, as well as potential hazards, to those who participate. Proponents of crowd work as a vehicle for global development highlight the potential to reduce the barriers to entry in the labor marketplace [52]. Participants are hired and paid based on tasks completed, as opposed to formal qualifications (such as degrees, employment history, or credible references) that are often out of reach in low-resource areas. The absence of a formal contract, and freedom from a physical workplace, allows anyone (in principle) to freely enter and exit the marketplace, and the flexibility of working hours enables part-time, supplemental earning as per the interest and availability of the worker.

On the other hand, there is increasing recognition that crowd work can also bring new hazards to workers [28]. Even if a crowd work platform is intended only for part-time, supplemental income, those who lack other secure sources of income (which encompasses most of the global poor) may come to work full-time, or even over-time, on such platforms. In this process, the end result of crowd work platforms could be to subvert intended labor regulations, for example, by paying less than minimum wage; by rewarding inhumane working hours; by enlisting help from underage workers; and in some cases, by issuing tasks that are of questionable moral or legal character, often unbeknown to the workers. Many of the safeguards associated with traditional employment, including associated benefits, avenues for appealing

decisions, and professional development, are usually absent from crowdsourced marketplaces. Analogous concerns have also been raised for the sharing economy [24].

In response to these important concerns, researchers have proposed various approaches and agendas for amplifying the benefits and equitability of crowd work platforms. Kittur et al. [37] lay out a compelling vision for a future of crowd work “in which we would want our children to participate”, while also highlighting the risks inherent in platforms as they exist today. Silberman et al. [49] propose high-level guidelines for treatment of crowdworkers, including paying at least the minimum wage at the worker’s location. Irani and Silberman [31] develop a tool, Turkopticon, to help protect workers from common hazards of Mechanical Turk, including low payments, arbitrary rejections of work, uncommunicative requesters, and pay delays.

We view our work as being in a similar spirit to those above, as we seek to improve the equitability of crowd platforms. The specific dimension we seek to impact is inclusiveness: enabling participation of low-income populations that were previously excluded. While such populations may have the most to benefit from crowd work, we acknowledge that they may also have the most to lose. There is no question that more effort is needed to ensure a beneficent future for the global marketplace of crowd work.

2.2 Crowdsourcing in Low-Income Communities

Our vision is to make digital work more accessible to NREGA workers in rural India. A majority of NREGA workers belong to the most disadvantaged sections of Indian society. In fact, a recent study noted that around 85 percent of NREGA beneficiaries belonged to Below Poverty Line families [41]. In short, a typical NREGA worker has very low income, is non-English speaking, and has low to non-existent levels of digital literacy [32]. While prior work on crowdsourcing has engaged with low-income communities, to our knowledge, this is the first work to examine the effectiveness of a crowdsourcing platform for digitization with NREGA workers from rural India. Introducing digital work to this population may also enable inclusion of workers who are not able to participate in NREGA, such as disabled individuals (unable to do manual labor) or women in certain communities (restricted mobility outside the home) [51].

mClerk [29] proposed a mobile crowdsourcing platform that enables digitization of handwritten Kannada text. The platform sends the image of a handwritten word as an SMS to the participants and expects them to respond with transliterated English text. Participants in the mClerk study were urban poor citizens from Bangalore who already owned a mobile phone and had proficiency in English. Our work builds on mClerk in two ways. First, all our participants are from rural India and were not proficient in English. Second,

since mClerk was published, the cost of smartphones has dramatically decreased. Our results show that native language input using a smartphone results in significantly higher accuracy (96.7%) compared to transliterated inputs as used by the mClerk platform (90.1% accuracy).

Similarly, MobileWorks [38] worked with urban poor in India to digitize English words and relied on a mobile web browser enabled by a data connection. This prevents MobileWorks from reaching our target demographic.

Perhaps the most prominent example of crowdsourcing in developing areas is Samasource [7]. Samasource offers digital work to its workforce via Internet-enabled computers set up in work centers or cybercafes. Samasource does not offer a mobile application for workers to work from home. In addition to other training, its workers undergo extensive training in speaking English as most of the provided work requires proficiency in English. Similarly, Gawade et al. [26] propose a platform that enables employment opportunities through microtasks via cybercafes. The study tests the English typing skills of participants using computers.

Our work differs from Samasource and Gawade et al. in two ways. First, while cybercafes can eliminate the need for the workers to own a computer, workers in rural India typically do not have access to such cybercafes. Secondly, most of our participants do not speak English and can only perform tasks in their local languages. More importantly, it is not immediately clear how efficiently one can digitize local-language text using a computer keyboard.

txtEagle, one of the earliest attempts in this space, gathered SMS-based survey responses from low-income people in Rwanda and Kenya [25]. Cellfare [14] makes the economic case for using microtasks as a form of social welfare. However, the work explored by Cellfare thus far was not from real-world tasks (e.g., reversing a 4-digit number). Like txtEagle, Cellfare delivers work over SMS messages.

Respeak [54] is a voice-based system that employs low-income users to transcribe audio files in English and Hindi. BSpeak [55] is a similar platform for blind users. Both systems worked with college students and rely on automatic speech recognition, which is still a work in progress for local dialects of Indian languages.

2.3 Other Crowdsourcing Platforms

In 2015, the Government of India launched the Digitize India Platform [40], a crowdsourcing platform that allows various government agencies to digitize their documents. Since its launch, the platform has been used to digitize 10.2 million documents [40]. Most of this work has gone to urban IT workers who have access to a computer and the Internet [40]. A major motivation for this paper is examining whether these handwritten-to-digital-text tasks can be conducted in rural Indian settings by the local population.

Platforms like Playment [5] and Captricity [2, 33] provide an end-to-end service to digitize documents. They automatically segment documents and use off-the-shelf Optical Character Recognition (OCR) to generate approximations. They then verify low-confidence OCR output using crowdsourced tasks. While our final platform may have a similar pipeline, there are two key differences. First, OCR for handwritten Indian languages is still not mature [30, 47, 50]. Second, these current platforms do not reach our target demographic.

There are several global crowdsourcing platforms such as CrowdFlower [53], Amazon Mechanical Turk [1], Cloud-Crowd, and others. Our work is distinguished from these efforts in several ways. First, it aims to work with rural Indians in resource-constrained areas with limited to non-existent digital literacy levels. Second, most of our participants cannot speak or comprehend English and they perform digital work exclusively in their local languages. Prior work has documented how language and usability challenges hinder usage of Mechanical Turk even among urban Indians [36]. Third, our participants usually do not have access to computers or broadband Internet.

3 SYSTEM DEVELOPMENT

We started by focusing on digitization of government documents in local languages. There are virtually unlimited volumes of paper documents in government offices, and there is also strong political will to convert these paper records to digital form [40]. However, the records are often handwritten and in local languages, making them unsuitable for automatic digitization using off-the-shelf OCR technology. A mobile crowdsourcing platform could leverage skill sets around regional languages to assist with this digitization. Such a platform would include the following steps:

- (1) manually scanning the documents
- (2) segmenting a scanned document into individual words.
- (3) automatically distributing them to multiple workers
- (4) making sure that the workers finish the task(s)
- (5) evaluating the accuracy of the submitted work
- (6) automatically aggregating the results to digital format

For this research paper, our focus is on steps 4 and 5. Specifically, our goal is to understand the accuracy with which rural workers can complete local-language digitization tasks. Although we expect the remaining steps to be similar to existing platforms (e.g., Captricity [2]), to make rural crowdsourcing platforms viable, we need to build low-cost mechanisms to reach and train rural workers. Given the proliferation of smartphones and growing access to cellular data, we are confident that future work can build scalable solutions to address this aspect of crowdsourcing.

3.1 Initial Application and Field Trial

We designed an Android application that allows us to measure the accuracy with which our participants can digitize handwritten words in their local language. To test the usability of our prototype, we conducted a two-week user study in Amale, a small tribal village in the Wada district of Western India [9]. The local language spoken in Wada and surrounding districts is Marathi (which follows the same Devanagari script as Hindi).

We worked with a local nonprofit based in Wada. In the last 2 years, the nonprofit has operated in several ‘*adivasi*’ (tribal) communities in the area. With their team, we visited the villages to gain an understanding of the existing digital literacy levels in the community, smartphone ownership, data penetration, etc.

3.2 Ethical Concerns

As our work targets vulnerable populations, sensitivity to research ethics is especially important. Dearden and colleagues [21, 22] provides an excellent overview of the unique ethical considerations in ICTD/HCI4D, as well as a call for increased vigilance in protecting local populations, especially when external researchers are engaged only temporarily in the community. In the case of our study (which was approved by an IRB), all field engagements were done collaboratively with long-term champions for the community: an organization in Amale [6] and a village leader in Soda [3]). These champions helped to ensure that the engagement was in the best interest of the community; that participants understood the engagement; and that there would be continuity of engagement to help the community benefit from any future outcomes/artifacts of the research. More specifically, champions conveyed that participation was completely voluntary; that participants could quit at any time; and that the study would run for only two weeks, as opposed to a permanent earning opportunity. As detailed later, our payment exceeded the local minimum wage. However, we did not position the research as an avenue for full-time employment or livelihood support. Instead, we sought to examine crowdwork holistically, both as supplemental income and for its potential to generate non-commercial engagement with digital content.

3.3 Participants

At the end of our first visit to Wada, we recruited a group of 12 people (8 men, 4 women) who volunteered to be a part of our first user study, after we announced the study in the village through our local partner. Every participant had completed at least a 5th standard education in their local language. The men in the group had completed at least 10th standard in a local government school. The group had people from the age of 18-29 and every participant belonged to a




Figure 1: Swarachakra Keyboard. Right image shows the chakra around the त character.

Below the Poverty Line family. Of the 12 participants, no one had access to a smartphone and only 2 people had access to a feature phone. For the duration of the study, we provided inexpensive Android smartphones that cost less than 50 USD to each of our participants. The tribal village of Amale also has no mobile connectivity. One of our participants informed us that once a day, every week, people who have feature phones climb a nearby mountain to get access to signal.

3.4 Swarachakra Keyboard

For our user studies, we used the Swarachakra keyboard [8, 35] designed by IIT-Bombay. Swarachakra is a logically structured, open-source keyboard for text input in Indic scripts. Swarachakra provides keyboards in both Hindi and Marathi, the languages we worked with for our studies. Most Indic language keyboards require many more keys, due to many matras (vowel modifiers) and sanyukts (compound words). The layout of the Swarachakra keyboard is based on the structure of the Indic scripts. Prior research has shown that doing so reduces the cognitive load that users often face while typing in Indic scripts compared to other alternatives (e.g., Google Indic Keyboard) and therefore, improves first-time usability [35]. As most of our users were using a smartphone for the first time, Swarachakra was a natural choice for our study. Figure 1 shows the Swarachakra keyboard.

3.5 Methodology

During our first visit to Amale, we conducted a few exercises with our participants where we mimicked real-world digitization tasks. These experiments were conducted on an Android smartphone we carried with us. We started by asking every participant to type their name on a standard Android Notes app in Marathi using the installed Marathi keyboard. We demonstrated the process to type your name on the phone once to the entire room. This training process took less than 5 minutes. Each participant took less than a minute to type their respective names. For all our participants, this was their first time typing on a smartphone.

It is common to see households in rural India sharing mobile phones. We modified the application to support multiple

accounts and this allowed more than one villager to use our app on the same phone.

As we could not find a large corpus of handwritten Marathi words, for the first study, we created a database to suit our needs. We computationally rendered 8000 Marathi words picked from the Swarachakra Marathi corpus [34] in a font that resembles handwritten Marathi. However, for our second study, we used a dataset of truly handwritten Hindi words (see Section 3.7).

A few months after our first visit, we visited Wada again with our updated prototype. Every participant was provided with the 8000 Marathi words they could digitize. At any time, the application would tell the participant how many words were still left undigitized.

On the first day of the study, a team of 3 researchers trained our 12 participants on how to use a smartphone and our application. We were able to use a room in the local government school (the only concrete structure in the village). Our participants were trained to use the smartphone by our team for 30 minutes on the first day. We attended to the participants individually and showed them how to locate our application on the phone and type words using the Swarachakra keyboard. There was no separate in-person training phase apart from the 30 minutes of training, and our participants learnt how to type Marathi on the Swarachakra keyboard while doing the work.

This part of the study was semi-structured. For the first 2 days, we worked with the participants and helped them with their queries. Most queries related to difficulty in typing certain words, which required a complicated combination of keys or sanyukts (compound words). In the Devanagari script, adjacent letters of compound words often modify each other. Sanyukts also need to be typed in one go, making them harder to type if one does not know the exact combination of keys. We also tried to have the participants help each other and when a participant would come to us asking for help with a word another participant had already digitized, we would direct them to the other participant. Participants were incentivized using bonuses and they were informed that the 3 participants with the highest accuracy rates would receive a bonus of INR 500 each at the end of the study.

During the study, we played several games, including a modified version of Telephone (where the first person would pick a random Marathi word from the application, whisper it in the ear of the person sitting next to them and so on, until the final person who would type the word they heard using the Swarachakra keyboard in the app).

We left the village after the first 2 days and the user study proceeded without any intervention from our side, for the next 10 days. On the 12th day, we came back to the village to observe the progress of our participants. The application stored all user performance data in a local database. We went

through the files for every participant and paid them in cash for their work. The mode of data collection had to be offline due to the lack of mobile signal in the village.

We sought to pay participants a rate that could be supported by potential future employers, yet higher than the existing wage offered by NREGA. In Maharashtra (the state where Amale is located), NREGA pays INR 201 per day [42]. Anticipating that participants could digitize at least 400 words per day, we decided to pay INR 0.5 per word, up to a maximum of 7200 paid words. This means that participants could earn INR 3600 for the two-week study (INR 300 a day for 12 days).

3.6 Outcome

Anecdotally, we observed that our participants could spend a minimum of 5-6 hours a day for our study and could digitize around 600 words per day. For our Amale study, our participants earned a minimum of INR 300 a day.

At the end of the 12-day study, every participant digitized over 8000 words and achieved an average accuracy of 94.5% in digitizing these words. We were able to measure accuracy easily as we already had the ground truth for every word and we simply compared the user input to the known ground truth. When asked, every participant stated that they felt the work was rewarding and fun and that they were fairly compensated for their work.

At the end of this study, we wanted to change three things with our next study. First, while we rendered handwritten words using a Marathi font, we wanted to perform a study with real handwritten words which are often poorly scanned and thus harder to comprehend. Second, most of our participants in Amale were young and unemployed. Thus, they were very excited to use a smartphone for the first time and have a job. We were interested in working with people across the age spectrum, with different job statuses and education levels. This is especially important as NREGA caters to a diverse set of rural Indians. Studies conducted in Rajasthan showed that most of the NREGA workers were found to be women and older men who had discontinued migration to the cities [48]. Third, some of our participants found the Swarachakra keyboard to be confusing and overwhelming initially. While every single participant gained a mastery over using the keyboard in a few days, we believe we can take steps to decrease the learning curve. To make the application more comfortable for first-time smartphone users, we wanted to improve our training program.

3.7 Building a More Accurate Database

The handwritten Hindi dataset that we used for our study in Soda is derived from the dataset created by Roy et al. [47]

for evaluating their new mechanism for recognizing handwritten Devanagari and Bangla scripts. Roy et al. have made the original dataset publicly available [12].

The Devanagari dataset contains 21334 images with each image containing a single word, with very few exceptions where the image had two words. While we would have ideally liked to use this dataset as is, we had to preprocess the dataset for our study for four reasons. First, different handwritten images had different amounts of white space surrounding them. This affects the size of the text when it is rendered by our application on the phone. Second, the label for each image is not available in Unicode, the current standard for representing text in languages other than English [20]. Third, not all images have the right label associated with them. Fourth, the dataset had images for only 2158 unique words, with multiple sets of handwritten images for each word.

To make the dataset suitable for our study, we first ran a segmentation algorithm that identifies an appropriate bounding box for the word in the image such that most of the white space around the word is removed [16]. We wrote a simple script that generates the Unicode label for each image from the label format provided in the original dataset [12]. After this process, we observed that several labels were incorrect with bad combination of Devanagari matras.

To fix the labels, we ran two rounds of verification. In the first round, each image was given to our team members who verified if 1) the image was cropped appropriately by our automated algorithm, and 2) if the label associated with the image is correct. If the label is incorrect, the team member can optionally type out the correct label. At the end of round one, the images that were not appropriately cropped or those that had incorrect labels were filtered out. In the second round, each image was again given to a team member to verify if the label is correct. At the end of the second round, we chose 6000 images with 1933 unique words such that each word was repeated at most 10 times. Note that each repeated instance is a different handwritten image of the same word.

4 USER STUDY IN SODA

For our second study, we choose Soda, a village in the state of Rajasthan. Soda falls in the Tonk district, which has been declared by the government as one of the 250 most resource constrained districts in the country [27]. The district is also one of the 12 districts in the state currently receiving funds from the Backward Regions Grant Fund Program [27]. We worked with the *sarpanch* (elected village head) of Soda to gain an understanding of the local communities.

4.1 The Soda Village

According to the 2011 Indian census, the Soda *panchayat* (administrative block) has over 5,000 people from different religions and nomadic communities with an average literacy

rate of 60.25%, lower than the national average of 74% [11]. The predominant language in Soda is Hindi. Being just over an hour from the state capital of Jaipur, the village has access to cellular data, and we observed that the smartphone penetration in the village was relatively high.

4.2 Participants

We worked with 32 people (25 men and 7 women) in the village who volunteered to be a part of our study. All participants had completed a 5th standard education in their local language and ranged in age from 18-38. The participants included unemployed college graduates, school teachers, homemakers, and college-going students. One of our participants had a known disability—he could not hear or speak. This diverse demographic allowed us to work with participants across the age spectrum, with different job statuses, and education levels.

We ensured that our user study did not detract our participants from their regular work (e.g., teaching) by clearly informing them that 1) they can do the work at any time of the day, and 2) there was no obligation on the part to complete all the words that we provided them.

Of the participants, 21 either already owned a smartphone or borrowed one from their father, brother, or husband (we observed that most smartphone owners were men). We provided the remaining 11 participants with a smartphone for the duration of the study. Participants who already owned a smartphone primarily used it for WhatsApp. In fact, some of our participants created a WhatsApp group to discuss queries on how to digitize certain words with each other.

4.3 Training

As mentioned before, the participants in the Amale study mentioned that they had difficulty in typing characters that combine multiple consonants (e.g., श्र pronounced shtra) in the Swarachakra keyboard. Therefore, we designed a more elaborate training for the Soda study. Based on the suggestions from the designers of the Swarachakra keyboard, we chose 42 Hindi words that covered most character combinations. We used these 42 words for training our participants to use the Swarachakra keyboard.

On Day 1 of the user study, we brought all our participants together in the Panchayat (village government) headquarters. We divided the 32 participants into three groups: 1) a group of young college-going students, 2) a group of government school teachers and their wives, and 3) a group of young unemployed men. We made sure that people in the same group either worked in the same workplace (e.g., a government school), lived close to each other or went to the same college. This was done to encourage members of a group to work with each other during the study.

Together, all 3 groups went through the 42 training words. We wrote the word on the blackboard and every group member typed the word in the application. The words progressively got more complicated and 3 researchers from our team helped participants who had any queries. At the end of the 1-hour training session, every participant had successfully digitized these 42 words.

Participants then started digitizing the actual words in their groups and once again, we helped them if/when they asked us a question. For some words, the spelling in the handwritten image was wrong compared to how the word is normally spelled. Many participants had queries regarding what to do in such cases. We asked our participants to type the corrected spelling. At the end of the first day, we asked one person per group to volunteer as the group leader. The group leader's primary job was to make sure that nobody in their team was falling behind in the work and to help each other, in case of any queries.

4.4 Study Protocol

In the study, all the words in our dataset were given to all the participants. We did this to measure the individual accuracy achieved by each participant. However, our eventual goal is to understand the feasibility of building a crowdsourcing platform with our participants. Therefore, to emulate a real-world crowdsourcing platform, we did not provide immediate feedback to our participants on whether each digitized word was correct/incorrect. Instead, we informed the participants of their performance at the end of week 1 and week 2.

The study lasted two weeks during which our participants were given 6000 words. On Day 1, we did the training with all our participants (as described above). On Day 2, we continued with the digitization exercise for 4 more hours and helped our participants with their questions. Once again, most questions were about the case where the spelling in the handwritten image was incorrect and queries about how to type certain compound words. By this time, most users would ask each other for help, instead of coming to us directly. We left the village at the end of Day 2. At the end of Week 2, we visited the village again to conduct interviews with our participants and distribute the payments for the work.

Based on our estimates from the Amale study, we told participants that we would pay INR 0.5 per correctly digitized word (though in practice, we paid INR 0.5 for all digitized words). With a total of 6000 words in the study, our participants could make up to INR 3000 over a period of 12 days, i.e., INR 250 per day on average. This is higher than both the NREGA minimum wage of INR 192 per day in Rajasthan [42] and the declared state minimum wage for semi-skilled work of INR 223 per day [13].

In addition to the payments mentioned above, we devised a bonus system to boost accuracy. The top 3 participants at the

end of the study (ranked on percentage of words accurately digitized) received a bonus of INR 1500 each, the next top 5 participants received a bonus of INR 1000 each, and the next top 10 participants received a bonus of INR 500 each. We observed that this encouraged a healthy competition between the participants.

5 QUANTITATIVE RESULTS

We analyzed the responses from all our participants for all the words. For each word, we identified the label provided by majority of the participants—we will refer to this as the majority response. Our analysis revealed that for 5727 out of the 6000 words, the majority response for each word matched the corresponding label in our database. While this result shows that majority of our participants already achieve good accuracy, we further analyzed the responses for the remaining 273 words for which the majority response did not match the label in the database. Our analysis of these 273 words revealed four categories of difficult word combinations.

First, we found that some characters in Devanagari can be rendered using multiple combinations of Unicode character set. When such characters are present in the word, the response from our participants may contain a different Unicode combination from the label in our database. In fact, we found that 209 of the 273 words had this mismatch. For example, consider the word आदि. The majority response was the Unicode combination आ + ढ + ि, while our label was the Unicode combination was अ + ि + ढ + ि. Both Unicode combinations result in the same word and are a result of typing the same word differently.

Second, for 17 words, either the spelling in the handwritten image was wrong compared to how the word is normally spelled or the handwritten image was ambiguous. In case of spelling errors, our participants had corrected the spelling. For ambiguous words, our interpretation agreed with that of the majority response. However, in these cases, the majority response did not match the database label.

Third, our database label was wrong for 44 words, whereas the majority response from our participants was correct. Finally, for just 3 out of the 273 words, the majority response from the participants was incorrect. Based on these results, we used the majority response from our participants as the correct response for our accuracy analysis. We ensured that this choice did not negatively affect our comparisons to a transcription firm.

5.1 Individual Performance

Figure 2 plots the overall accuracy achieved by each of our 32 participants. Our results show that our participants achieved an accuracy of 96.7% on average, with every participant




Figure 2: Individual participant performance. Triangle markers correspond to women users and red markers correspond to first-time smartphone users.

achieving a minimum accuracy of 93%. More importantly, 21 of the 32 participants achieved accuracy more than 96%.

The figure also highlights the accuracy for women users (triangle markers) and first-time smartphone users (red markers). Women’s participation in NREGA varies depending on their care responsibilities that limit their mobility outside their homes and available time for paid work [45]. Digital work can potentially provide more flexible work to women. Given the lack of access to smartphones, we expected men to perform better than women. Similarly, we expected regular smartphone users to perform better than first-time users. While our results match these expectations, our analysis indicates that the difference in accuracy, while statistically significant, is just 2% for both comparisons.¹

5.2 Crowdsourcing Platform

Since we had given all the 6000 words to all participants, we were able to simulate the effect of different crowdsourcing algorithms using the responses from our participants. For instance, in our simplest crowdsourcing algorithm, we would provide each word to two randomly chosen participants. If the responses from both our participants matched, then we would consider that as the crowdsourced response. If not, we would iteratively provide the word additional participants until two participants provide the same response, which will be considered the crowdsourced response. To improve the accuracy, we can increase the number of matching responses from 2 to higher values. Note that this improved accuracy will also come with increased cost per word. We simulated these crowdsourced algorithms on our data.

Figure 3 shows the results of this simulation for different values of matching response required by the crowdsourcing platform. The leftmost point corresponds to the case where each word is given to only one participant. In this case, the platform would achieve an accuracy of 96.7% at an amortized cost of INR 0.58 per word (INR 0.5 base payment plus INR

¹Male/Female: $t(30) = 2.22, p < 0.05$; Phone expertise: $t(30) = 2.07, p = 0.05$.




Figure 3: Cost vs. accuracy for 1) crowdsourcing platform with our participants (x), and 2) a transcription firm (■).

0.08 due to the bonuses offered). We draw two other conclusions from our results. First, even when requiring only two matching responses (the second leftmost “x” point), the crowdsourcing platform improves the accuracy to 98.9% at a cost of INR 1.19 per word. Second, as expected, increasing the number of matching responses to higher values (further “x” points on the line) increases the accuracy albeit at an increased cost. However, the accuracy plateaus at 99.8%.

5.3 Comparison to a Transcription Firm

To compare the performance of our participants to existing market alternatives, we contacted two transcription firms that offered digitizing handwritten text as a service. The cost charged by the two firms to digitize one word was INR 9.00 and INR 1.00. Since INR 1.00 per word was closer to the cost of our model, we gave the 6000 words to that transcription firm. To ensure that we were fair in our comparison, we assumed that a response from the transcription firm was correct if it matched either the label in our database or the majority response from our participants. After receiving the results, we noticed that their responses had Unicode mismatch errors similar to our platform. Once we accounted for these cases, the accuracy of the firm was 93.5%. Further analysis of the responses revealed many systemic errors. For instance, most responses were missing the nuqta (dot below a character, e.g., ढ vs. ढ̣). Assuming this error can be fixed, the accuracy goes up to 96.8%. Even if we assume all systemic errors can be fixed, the accuracy of the firm was 98.4%.

Figure 3 plots these points in comparison to our crowdsourcing model. Our simplest model that waits for two matching responses achieves a better accuracy (98.9%) at a slightly higher cost (INR 1.19). These results validate the feasibility of a crowdsourcing model using workers from our target demographic. While our study only used words from the Devanagari script, we believe our results will extend to other regional languages. More importantly, we believe a crowdsourcing platform using workers from rural India can be a

viable method to perform digital work in regional languages for which there are no existing options in the market.

6 QUALITATIVE RESULTS

At the end of the Soda user study, we conducted interviews with all our participants individually. The interviews were conducted by members of the research team in a local government building. Cognizant of the potential bias of this approach, being both outsiders and conducting our work at an official location, we attempted to put the participants at ease by presenting the interviews as conversations and encouraging all forms of feedback. We recorded and later, transcribed these interviews in Hindi. The interviews were hand-coded for themes by the primary author of the paper, following which quotes highlighting the key themes were curated to provide descriptive detail of respondents’ attitudes towards our work.

6.1 Work Flexibility and Increased Income

Our participants largely reported various positive attitudes towards the prospect of digital work. They articulated a range of reasons for the same, including the potential financial benefit from new work, and the pleasure of being able to participate in a digital project. The interactions on the system were novel to the participants. Many had never entered text in the Devanagari script before, and respondents reported enjoying the tasks for the pleasure of using their native language.

I really like this kind of work because there’s no physical pain or hard work involved. Because we are villagers, most of the work we get requires heavy lifting. But now, I can work from my home. I like that it matches my routine. I also really liked typing in shudh Hindi (classical Hindi).

—P13, 25 years old, male, village government helper

Several participants reported preferring this form of work over the physical work distributed by NREGA, which could be rigid in terms of workplace location, and involved a significant time and energy investment. NREGA only provides a limited amount of work annually and it can be inflexible regarding periods of work available. So, for individuals who were unemployed or underemployed, the prospect of an additional source of income was important.

I currently don’t have a job. I really benefited a lot due to this work. I was able to pay the bills for my home. I can use this money to pay for all the things I need this month, I don’t need to borrow money from anyone else. This work takes less time and energy than manual work, and the pay is much better.

—P19, 19 years old, male, unemployed

Some of these responses were indeed driven by the one-time experience of working on our system, whereby our

participants were unsure of what to expect in the future in terms of potential wage rate or employable hours. Nonetheless, they showed optimism about the opportunities this opened because of their existing underemployment. One way in which at-home work opened peoples' schedule was allowing them to use their free time more productively. As one user put it:

I think this work can really help our village. People who are unemployed right now will get some work and they will get to learn to type in their own language, so it's very good. Earlier, they would sit together and gossip. But, now they have work to do, so less time is wasted.

—P12, 24 years old, male, government school teacher

For the vast majority of our respondents, work typically involved physical labor, or work in traditional sectors of the economy (e.g. farming). Thus, specific parts of the day were often necessarily blocked off, and other parts of the day were invariably free. The flexible work schedules were helpful.

I would wake up at 4am every day and start digitizing words, as I had to leave for work at 7 am. Then, all the teachers doing this work would meet during lunch time in the school, sit together in one group and digitize together. In the evening, I have to take the buffaloes to the pond and while they were bathing and drinking water, I would use do this work.

—P12, 24 years old, male, government school teacher

The quote above also highlights another trend we noticed repeatedly — extreme enthusiasm to complete the work. The lack of workplace opportunities in the economy meant that our participants were attached to the possibilities that such work offered. In one case, a participant pulled two consecutive all-nighters to digitize all 6000 words within the first three days of receiving the work. This participant was 23 years old and could neither hear nor speak; this was the first employment of his life. We were obviously very concerned by this extreme working pattern, and discuss how it might be prevented in Section 7.

I really enjoyed doing this work. There's no time wasted in my day anymore. Because of school vacations, I have nothing to do right now. So, instead of giving me 3000 words in a week, if you give me 10,000 words next time, my free time will be better used. I will make more money and my life will be fixed.

—P13, 25 years old, male, village government helper

While the respondent's performance of the tasks, as well as his sense of self-worth related to work were very positive, his aspiration of having the work enable a livelihood underlined the risks of expectations that such work may build up for people with limited employment prospects.

6.2 Devanagari Interface Comfort

As we mentioned previously, typing in Devanagari script can be challenging. Most words need multiple keystrokes and modifier combinations. A few of our participants who owned smartphones told us that even though they couldn't speak English, they preferred typing Hindi in the transliterated Romanized script using an English keyboard before this user study. Prior research has shown that Romanized text imposes a significant neurocognitive load on its readers [44] and most participants who typed in Romanized script told us that they simply "didn't type much on their smartphones". Thus, for most of our participants, this was their first time typing in Hindi (or Marathi) on a smartphone. Participants self-reported comfort with typing using Devanagari script. For those who had prior experience with transliterative typing, this was an opportunity to type in shudh ("non-anglicized", in this context) Hindi. As one of our participants put:

My typing speed improved, and I learnt how to type in shudh Hindi. Earlier, I didn't type much on my phone as my typing speed was so slow.

—P1, 19 years old, female, college student

As prior research has shown, the first few instances of using a smartphone interface can be crucial in convincing people about their ability to use certain technology [19]. Most participants reported increased confidence with being able to use a smartphone after the study. For some participants, this was their first time doing paid work. We found that despite what we may classify as initial successes (basic task completion), respondents took a few days before feeling comfortable enough with their ability to independently manage the tasks.

I now feel that I am capable of doing any new work. I didn't feel like that before. First 2-3 days, I wasn't sure if I will be able to do this work. But slowly, I figured it out. If you give me some different new work today, I think I can do it.

—P10, 38 years old, male, government school teacher

Typing in a local language, and completing tasks offered a departure from notions that doing complex technical things was somehow a domain of elite populations. Some of our participants reported a sense of linguistic pride in being able to type fluently in their own language. Several users reported starting using Hindi in messaging apps like WhatsApp as an outcome of their involvement in the study. As one user said:

After doing this work, I type a lot more in Hindi. Look, we even created a Whatsapp group to discuss village problems and I have replied in Hindi. It is really nice that we all learnt to type in Hindi. It is a great language, plus it's our language. So everyone should know how to type in it.

—P10, 38 years old, male, government school teacher

6.3 Convivial Sensemaking

We observed that digitizing words became a communal activity, in part because of the unfamiliarity with the interface, but more importantly because of the discursive nature of the activity itself. The tasks involved checking for techniques on input, confirming doubts about characters read, or simply working together in a convivial, playful nature. People digitized words together with both friends and family members.

Whenever I had a doubt on how to digitize a word, I would message my friends and we would go to school early and digitize together. Since I was the group leader, I would make sure they were all doing okay.

—P13, 25 years old, male, village government helper

Whenever I had a doubt, I would ask my sister and family for help. When I got bored, my father and brother digitized some words on my behalf, because they were excited to do this work.

—P1, 19 years old, female, college student

The notion of convivial design implies a playful, ludic style of building, which designs up from what participants find as fun, and engineers that into the interactions of the artifact. While ludic design approaches were not explicitly integrated into the processes we followed in building our application, elements of playful exchange and even design discussion about what could be optimized in the interface were clear in the discussions that emerged.

My sister is also doing this work, and we have a competition going on. Last week, she got 92% and I got 91%,² but I am going to make sure that I beat her in the end.

—P8, 22 years old, female, unemployed

The weaving of crowd tasks into social interactions was also a reflection of daily life in Soda. Convivial principles already existed in the social lives of our respondents; especially agriculture, which is communal in nature. People in Soda lived in close quarters, knew each other and often worked together.

I live on the same street as two other people who were doing this work. So, we would all meet at my house every day and digitize together, so it was more fun.

—P4, 18 years old, male, college student

The playful nature of the device use was sometimes explicitly verbalized in terms of breaking the monotony of life as usual. With limitations on the kinds of wage work or educational options available, few retail or recreational facilities, the crowdsourcing tasks blended into leisure.

I would just sit idle for most of the day before this work. Whenever I didn't feel like doing college work, I would open the app

²Our initial accuracy feedback was based on comparisons to the label in the database and not the majority response. Hence, the participant reported lower numbers.

and start digitizing. And I wouldn't get bored at all. I thought of it like a game.

—P1, 19 years old, female, college student

There was some tension in how the devices were understood as personal objects. The notion of a cell phone as something more than a device exclusively for “play” emerged in discussions as respondents pointed out the initial surprise that something considered fun to work on can also lead to productive economic activity.

First, my parents scolded me for spending so much time on the phone. Once I told them that I was doing work and getting paid, they scolded me when I wasn't doing work on my phone.

—P8, 22 years old, female, unemployed

Living in precarity, access to this work also meant that some of the same respondents who recognized the playful aspects of this work, nevertheless saw this work as serious business and were careful about making sure the work was done right.

I didn't let anybody else do work on my behalf. What if they get it wrong? Everyone told me that my bhabhi (brother's wife) is educated and she can do the work with me. But, I didn't let her. I wanted to get full 100% accuracy.

—P8, 22 years old, female, unemployed

7 DISCUSSION

On the whole, our findings could be interpreted as a positive indicator for the viability of crowdsourcing work to rural areas. Participants were eager to engage with supplemental work. They were able to perform digitization tasks accurately—rivaling that of leading market alternatives—despite limited digital literacy, including lack of prior experience with local language text entry. In addition, they were able to complete tasks quickly enough that they could earn more than existing local wages, while charging similar rates as other transcription firms. All of these factors point to a potential future for crowdsourced work as a means of supplemental income in this demographic.

At the same time, we also found examples where digital work disrupted daily life, for example, by enabling unacceptably long working hours. While some could argue that individuals should be empowered with the agency to choose their own hours, others could view the choices made as self-exploitation that is a violation of humane labor laws. Our design recommendation for future platforms is to include frequent nudges for taking breaks, and to impose a hard limit of working hours after a certain point. Such limits are consistent with intended use as a supplemental source of income, as opposed to a full-time livelihood.

Our study has some limitations. First, for many of our participants, this study was their first time using a smartphone.

Our study lasted for 12 days. While most of our participants took the full duration to complete the work given to them, 12 days might not have been enough for the novelty factor to wear off. On the other hand, with a longer duration, participants might acquire more skills and become more efficient in performing the given tasks. Second, since the participant interviews were conducted by the researchers themselves, there is a potential response bias in our interviews [23]. To mitigate this factor, we 1) presented the user interviews as friendly conversations to put our participants at ease, 2) asked open-ended questions and avoided leading questions, and 3) encouraged all forms of feedback.

In our user studies, our participants were trained in person by one of the researchers. However, such in-person training might not be scalable to a large number of workers. As part of our future work, we would like to explore the effectiveness of remote training based on videos. Given the communal nature of life in rural communities, we would also like to explore the use of collaborative training with incentives.

The digital work provided by our study assumes that our participants can read and write a local language (Hindi and Marathi, in our case). However, we cannot assume that all NREGA beneficiaries are literate. Literacy rates among NREGA workers vary from state to state, from 20% literacy in parts of Rajasthan [15, 45] to 80% literacy in Kerala [4]. While our results demonstrate that low-income workers in rural India can accurately digitize local-language documents, to build a truly inclusive model, we need to experiment with forms of digital work that do not require literacy.

Several participants told us that they thought of the work as a game. Some participants said that when they got bored of the work, they would listen to music on their phones while digitizing words. Such anecdotes and insights suggest that we should implement elements of playful design. Past work has proposed that ICTD can benefit from ludic, contextual design [18]. During our studies, we noticed friendly competition among our participants to achieve higher accuracy rates. As we described, most participants in the village worked with each other. We would like to experiment with features such as an in-village leaderboard to strengthen these factors.

8 CONCLUSION

This paper makes a case for creating a crowdsourcing platform that employs rural Indians to complete digital work that involves tasks in local languages. These are the kinds of work that the Indian government seeks to do with its Digital India initiative. We conducted user studies in two resource-constrained villages in India. For most of our participants, this was their first time using a smartphone. Our participants outperformed an existing transcription service in terms of accuracy of digitizing handwritten Hindi text for a similar cost. While our research does not answer questions

of whether digital work can be a viable, fulfilling means of primary livelihood, we show that there is potential for supplemental work and demonstrate the economic viability of building a crowdsourcing platform that employs our target demographic.

ACKNOWLEDGEMENTS

We thank the reviewers of our paper for their valuable comments on our paper. We are grateful to Anuja and Pranita from IIT Bombay, Rohit and Nikesh from Rural Caravan, Chhavi Rajawat from Soda and Tanuja from MSR India for their immense help and support during our field studies.

REFERENCES

- [1] Amazon Mechanical Turk. www.mturk.com.
- [2] Captricity. <https://captricity.com/>.
- [3] Chhavi Rajawat. https://en.wikipedia.org/Chhavi_Rajawat/.
- [4] "MNREGA in Kerala: Summary of Findings, Suggestions and Conclusion. http://shodhganga.inflibnet.ac.in/bitstream/10603/49084/18/18_summary_conclusion.pdf.
- [5] Playment: Training Data for AI, Data Enrichment Services. <https://playment.io/>.
- [6] Rural Caravan. <http://www.ruralcaravan.com/>.
- [7] Samasource. <https://www.samasource.org/>.
- [8] Swarachakra. <http://idid.in>.
- [9] 2011. Amale Village. <https://www.censusindia.co.in/villages/amale-population-thane-maharashtra-552004>.
- [10] 2011. Socio-Economic Caste Census. <http://secc.gov.in/categorywiseIncomeSlabReport?reportType=All%20Category>.
- [11] 2011. Soda Village. <https://www.census2011.co.in/data/village/92572-sodarajasthan.html>.
- [12] 2016. Indic Word Dataset. <http://www.iitr.ac.in/media/facspace/proy.fcs/IndicWord.rar>.
- [13] 2018. Latest Minimum Wage in Rajasthan. <https://paycheck.in/salary/minimumwages/rajasthan>.
- [14] Erlend Berg, D Rajasekhar, and R Manjula. 2017. *Cellfare: Delivering Self-Targeted Welfare Using Mobile Phones*. CSAE Working Paper Series. Centre for the Study of African Economies, University of Oxford. <https://ideas.repec.org/p/csa/wpaper/2017-14.html>
- [15] Rhonda Breitreuz, Carley-Jane Stanton, Nurmaiya Brady, John Pattison-Williams, ED King, Chudhury Mishra, and Brent Swallow. 2017. The Mahatma Gandhi National Rural Employment Guarantee Scheme: A Policy Solution to Rural Poverty in India? *Development Policy Review* 35, 3 (2017), 397–417.
- [16] Thomas M Breuel. 2008. The OCRopus Open Source OCR System. In *Document Recognition and Retrieval XV*, Vol. 6815. International Society for Optics and Photonics, 68150F.
- [17] William Bright. 1996. The Devanagari Script. *The World's Writing Systems* (1996), 384–390.
- [18] Padma Chirumamilla and Joyojeet Pal. 2013. Play and Power: A Ludic Design Proposal for ICTD. In *Proceedings of the Sixth International Conference on Information and Communication Technologies and Development: Full Papers - Volume 1*.
- [19] Deborah R. Compeau and Christopher A. Higgins. 1995. Computer Self-Efficacy: Development of a Measure and Initial Test. *MIS Quarterly* 19, 2 (1995), 189–211.
- [20] Unicode Consortium et al. 1997. *The Unicode Standard, Version 2.0*. Addison-Wesley Longman Publishing Co., Inc.

- [21] Andy Dearden. 2012. See No Evil?: Ethics in an Interventionist ICTD. In *Proceedings of the Fifth International Conference on Information and Communication Technologies and Development*.
- [22] Andy Dearden and Dorothea Kleine. 2019. Ethical standards for the ICTD/ICT4D community: A participatory process and a co-created document. In *Proceedings of the Tenth International Conference on Information and Communication Technologies and Development*.
- [23] Nicola Dell, Vidya Vaidyanathan, Indrani Medhi, Edward Cutrell, and William Thies. 2012. "Yours is Better!": Participant Response Bias in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
- [24] Tawanna Dillahunt, Airi Lampinen, Jacki O'Neill, Loren Terveen, and Cory Kendrick. 2016. Does the Sharing Economy Do Any Good?. In *Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion*.
- [25] Nathan Eagle. 2009. Tختهagle: Mobile Crowdsourcing. In *Proceedings of the 3rd International Conference on Internationalization, Design and Global Development: Held As Part of HCI International 2009*.
- [26] Mrunal Gawade, Rajan Vaish, Mercy Nduta, and James Davis. 2012. Exploring Microwork Opportunities Through Cybercafés. In *Proceedings of the Second ACM Symposium on Computing for Development*.
- [27] Ruchi Ghose and Vinod Jain. 2016. *Starting and Sustaining Voluntary Teacher Forums: Experience from Tonk, Rajasthan*. Technical Report.
- [28] Mary L. Gray and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Eamon Dolan/Houghton Mifflin Harcourt.
- [29] Aakar Gupta, William Thies, Edward Cutrell, and Ravin Balakrishnan. 2012. mClerk: Enabling Mobile Crowdsourcing in Developing Regions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
- [30] M. Hanmandlu, O. V. Ramana Murthy, and Vamsi Krishna Madasu. 2007. Fuzzy Model Based Recognition of Handwritten Hindi Characters. In *Proceedings of the 9th Biennial Conference of the Australian Pattern Recognition Society on Digital Image Computing Techniques and Applications*.
- [31] Lilly C. Irani and M. Six Silberman. 2016. Stories We Tell About Labor: Turkopticon and the Trouble with "Design". In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*.
- [32] Nimisha Jain and Kanika Sanghi. 2016. The Rising Connected Consumer in Rural India. http://image-src.bcg.com/images/BCG-The-Rising-Connected-Consumer-in-Rural-India-July-2016_tcm9-61868.pdf.
- [33] Huguens Jean, Yoriyasu Yano, Hui Peng Hu, and Kuang Chen. 2017. Analyzing content of digital images. US Patent 9,652,688.
- [34] Anirudha Joshi, Girish Dalvi, and Manjiri Joshi. 2014. Corpus of Marathi Word Frequencies from Touch-Screen Devices Using Swarachakra Android Keyboard. In *Proceedings of the India HCI 2014 Conference on Human Computer Interaction*.
- [35] Anirudha Joshi, Girish Dalvi, Manjiri Joshi, Prasad Rashinkar, and Aniket Sarangdhar. 2011. Design and Evaluation of Devanagari Virtual Keyboards for Touch Screen Mobile Phones. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*.
- [36] Shashank Khanna, Aishwarya Ratan, James Davis, and William Thies. 2010. Evaluating and Improving the Usability of Mechanical Turk for Low-income Workers in India. In *Proceedings of the First ACM Symposium on Computing for Development*.
- [37] Aniket Kittur, Jeffrey V. Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. 2013. The Future of Crowd Work. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*.
- [38] Prayag Narula, Philipp Gutheim, David Rolnitzky, Anand Kulkarni, and Bjoern Hartmann. 2011. MobileWorks: A Mobile Crowdsourcing Platform for Workers at the Bottom of the Pyramid. In *Proceedings of the 11th AAAI Conference on Human Computation*.
- [39] Department of Electronics & Information Technology Government of India. 2015. Digital India Programme. <http://www.digitalindia.gov.in>.
- [40] Department of Electronics & Information Technology Government of India. 2015. Digitize India Platform :: Transforming Pixels to Data. <https://digitizeindia.gov.in/>.
- [41] Ministry of Rural Development Government of India. MGN-REGA Sameeksha, An Anthology of Research Studies on the Mahatma Gandhi National Rural Employment Guarantee Act, 2005, 2006-2012. http://nrega.nic.in/Circular_Archive/archive/MGNREGA_SAMEEKSHA.pdf.
- [42] Ministry of Rural Development Government of India. NREGA Wage Rate. http://nrega.nic.in/netnrega/writereaddata/Circulars/2058Notification_wage_rate_2017-2018.pdf.
- [43] Ministry of Rural Development Government of India. 2005. The Mahatma Gandhi National Rural Employment Guarantee Act. <http://www.nrega.nic.in>.
- [44] Chaitra Rao, Avantika Mathur, and Nandini C Singh. 2013. 'Cost in Transliteration': The Neurocognitive Processing of Romanized Writing. *Brain and language* 124, 3 (2013), 205–212.
- [45] Savitri Ray and Madhuri Karak. Women in NREGA: Issues of Child Care. <http://www.forces.org.in/publications/NREGA%20Report.pdf>.
- [46] Joel Ross, Lilly Irani, M. Six Silberman, Andrew Zaldivar, and Bill Tomlinson. 2010. Who Are the Crowdworkers?: Shifting Demographics in Mechanical Turk. In *CHI '10 Extended Abstracts on Human Factors in Computing Systems*. 2863–2872.
- [47] Partha Pratim Roy, Ayan Kumar Bhunia, Ayan Das, Prasenjit Dey, and Umapada Pal. 2016. HMM-based Indic Handwritten Word Recognition Using Zone Segmentation. *Pattern Recogn.* 60, C (Dec. 2016), 1057–1075.
- [48] Tushaar Shah, Shilp Verma, R Indu, and P Hemant. 2010. Asset Creation Through Employment Guarantee: Synthesis of Student Case Studies in 9 states of India. *Anand: International Water Management Institute (IWTMI)* (2010).
- [49] M. S. Silberman, B. Tomlinson, R. LaPlante, J. Ross, L. Irani, and A. Zaldivar. 2018. Responsible Research with Crowds: Pay Crowdworkers at Least Minimum Wage. *Commun. ACM* 61, 3 (Feb. 2018), 39–41.
- [50] Gunjan Singh and Sushma Lehri. 2012. Recognition of Handwritten Hindi Characters Using Backpropagation Neural Network. *International Journal of Computer Science and Information Technologies* 3, 4 (2012), 4892–4895.
- [51] Ratna M Sudarshan. 2011. *India's National Rural Employment Guarantee Act: Women's Participation and Impacts in Himachal Pradesh, Kerala and Rajasthan*. Technical Report.
- [52] William Thies, Aishwarya Ratan, and James Davis. 2011. Paid Crowdsourcing as a Vehicle for Global Development. In *CHI Workshop on Crowdsourcing and Human Computation*.
- [53] Chris Van Pelt and Alex Sorokin. 2012. Designing a Scalable Crowdsourcing Platform. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*.
- [54] Aditya Vashistha, Pooja Sethi, and Richard Anderson. 2017. Respeak: A Voice-based, Crowd-powered Speech Transcription System. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*.
- [55] Aditya Vashistha, Pooja Sethi, and Richard Anderson. 2018. BSpeak: An Accessible Voice-based Crowdsourcing Marketplace for Low-Income Blind People. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*.