

Crowdsourcing the Installation and Maintenance of Indoor Localization Infrastructure to Support Blind Navigation

COLE GLEASON, Carnegie Mellon University, USA

DRAGAN AHMETOVIC, Carnegie Mellon University, USA

SAIPH SAVAGE, West Virginia University, USA

CARLOS TOXTLI, West Virginia University, USA

CARL POSTHUMA, National Autonomus University of Mexico, Mexico

CHIEKO ASAKAWA, Carnegie Mellon University, USA

KRIS M. KITANI, Carnegie Mellon University, USA

JEFFREY P. BIGHAM, Carnegie Mellon University, USA

Indoor navigation systems can make unfamiliar buildings more accessible for people with vision impairments, but their adoption is hampered by the effort of installing infrastructure and maintaining it over time. Most solutions in this space require augmenting the environment with add-ons, such as Bluetooth beacons. Installing and calibrating such infrastructure requires time and expertise. Once installed, localization accuracy often degrades over time as batteries die, beacons go missing, or otherwise stop working. Even localization systems installed by experts can become unreliable weeks, months, or years after the installation. To address this problem, we created LuzDeploy: a physical crowdsourcing system that organizes non-experts for the installation and long-term maintenance of a Bluetooth-based navigation system. LuzDeploy simplifies the tasks required to install and maintain the localization infrastructure, thus making a crowdsourcing approach feasible for non-experts. We report on a field deployment where 127 participants installed and maintained a blind navigation system over several months in a 7-story building, completing 455 tasks in total. We compare the accuracy of the system installed by participants to an installation completed by experts with specialized equipment. LuzDeploy aims to improve the sustainability of indoor navigation systems to encourage widespread adoption outside of research settings.

CCS Concepts: • **Human-centered computing** → **Computer supported cooperative work**; **Accessibility technologies**; Ubiquitous and mobile computing systems and tools; • **Hardware** → **Sensor applications and deployments**; • **Computer systems organization** → *Sensor networks*; • **Social and professional topics** → *People with disabilities*;

Additional Key Words and Phrases: Individuals with visual impairments; Real-world accessibility; Physical crowdsourcing; Indoor navigation assistance; Volunteer crowd work.

This work was sponsored in part by JST CREST grant (JPMJCR14E1), NSF NRI grant (1637927).

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. (NSF grant number). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation."

Authors' addresses: Cole Gleason, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213, USA; Dragan Ahmetovic, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213, USA; Saiph Savage, West Virginia University, 401 Evansdale Drive, Morgantown, WV, 26506, USA; Carlos Toxtli, West Virginia University, 401 Evansdale Drive, Morgantown, WV, 26506, USA; Carl Posthuma, National Autonomus University of Mexico, Av Universidad 3000, Ciudad Universitaria, Mexico City, Distrito Federal, 04510, Mexico; Chieko Asakawa, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213, USA; Kris M. Kitani, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213, USA; Jeffrey P. Bigham, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213, USA.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2018 Copyright held by the owner/author(s).

2474-9567/2018/3-ART0

<https://doi.org/0000001.0000001>




Fig. 1. Examples of LuzDeploy tasks: (a) The bot directs a worker to the general area in which a beacon should be placed; (b) a subsequent sampling step (also performed by non-experts) creates a model of the Bluetooth signals used for localization.

ACM Reference Format:

Cole Gleason, Dragan Ahmetovic, Saiph Savage, Carlos Toxtli, Carl Posthuma, Chieko Asakawa, Kris M. Kitani, and Jeffrey P. Bigham. 2018. Crowdsourcing the Installation and Maintenance of Indoor Localization Infrastructure to Support Blind Navigation. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 0, Article 0 (March 2018), 25 pages. <https://doi.org/0000001.0000001>

1 INTRODUCTION

People with vision impairments can have difficulty navigating unfamiliar indoor spaces because visual cues and signage are not accessible to them [55]. The dominant approaches to make indoor environments accessible to people with visual impairments require installing infrastructure add-ons [18, 37, 43, 56], but this is often prohibitively expensive to do at scale and maintain over time in the real world. Typically, experts must install the localization system and perform on-going maintenance to ensure a high level of accuracy. Installation can be cost prohibitive and on-going maintenance is often too infrequent to ensure infrastructure accuracy when a person arrives who needs to use it.

In this paper, we introduce LuzDeploy, a system that coordinates a local crowd to install and maintain an indoor localization system via a Facebook bot using supporting mobile applications designed for non-experts. We focus on indoor navigation systems that use Bluetooth Low Energy (BLE) beacons placed densely in the environment, as this method is capable of achieving sub-meter localization accuracy [1]. Because this approach uses commodity hardware, a person with visual impairments just needs a smartphone, while building owners only need to place relatively inexpensive BLE beacons. However, installing the beacons takes time and expertise. Furthermore, as beacons go missing or batteries die the localization accuracy degrades over time, so a trained administrator must continually return to maintain the system.

Installing navigation infrastructure is currently an immense, upfront, and complex process performed by people with expertise. Maintenance over the weeks, month, and years following installation is rarely done frequently enough, and so localization accuracy naturally degrades as the infrastructure ages. When someone with visual impairments enters a building, how can they trust that their navigation system will work? The long-term sustainability of these systems hinge on experts maintaining the infrastructure over years, but because there are few people with adequate training this does not scale well to large real-world environments. Administrators

might hire and train more experts, but this also increases the cost of indoor localization technology. Because the hiring and training process takes time to complete, it delays immediate adoption. Even when recruiting a large workforce is possible (*e.g.*, a pool of volunteers), coordinating them for concurrent installation or continuous maintenance is challenging and time-consuming. In particular, to ensure the accuracy of the system, the workers need to quickly notice and repair sporadic beacon failures.

LuzDeploy is inspired by recent work on crowdsourcing and collective action demonstrating that people can be recruited via an open call and coordinated to work together to complete substantial tasks that have traditionally required experts [6]. Such approaches are feasible because workers complete relatively small, decontextualized tasks, while the system coordinates their work by planning tasks and maintaining higher-level context about what needs to be done. In LuzDeploy, crowd members place and replace beacons, check the health of existing beacons, and sample Bluetooth signals in particular locations to help the system learn to detect these locations in the future. At no point does the crowd member need to understand how indoor localization works, but rather only needs to perform simple tasks that we show do not require expert training.

LuzDeploy has two important features to engage the non-expert mobile crowd:

- 1 It allows easy *non-expert* participation: LuzDeploy simplifies a complex installation and maintenance process by distributing small tasks with step-by-step, just-in-time instructions. A worker can quickly complete small tasks with no expert training by following the given steps.
- 2 It enables *drop-in, anytime* engagement: The previously long process of instrumenting an entire building is broken into small, predictable batches of work that workers can complete in minutes. Instead of performing maintenance all at once every few months, the process is also spread out over these small tasks. This allows passerby with just a few minutes of availability to quickly sign up and significantly impact the collective effort.

We investigate how the installation and maintenance of such systems can be completed by non-expert workers. We applied LuzDeploy to install and maintain a blind navigation system, NavCog [1], in 7 floors of the computer science building at Carnegie Mellon University. Overall, 127 participants used LuzDeploy to install BLE beacons and sample signal data throughout the building. This was completed in four different sessions:

- (1) First session: We used LuzDeploy to coordinate a large number of drop-in, non-expert volunteers during beacon installation. Over the course of one afternoon, 89 workers quickly joined the installation effort and placed 99 BLE beacons in the building.
- (2) Second session: To simulate the slow and sporadic nature of infrastructure maintenance tasks, a second beacon deployment session took place over three months, during which 25 participants were engaged casually based on their availability and placed 97 beacons.
- (3) Third session: To see if we could encourage more involvement in data sampling tasks, we provided small amounts of compensation for the participants. Over two weeks, nine participants completed 25 data-sampling tasks and collected 185 Bluetooth signal samples at 112 unique locations in the building.
- (4) Fourth session: Our third session did not give us enough information about the feasibility of crowdsourcing to collect signal samples, so we held an additional month long session targeting a single floor of the building. During this period, eight participants completed 211 tasks, collecting 1638 Bluetooth signal samples at 116 unique locations.

Our work shows generally that it is possible to engage non-expert participants in effective physical crowdsourcing, and shows specifically that we can effectively use this approach to install and maintain indoor localization infrastructure. Additionally, by automating the long-term maintenance of NavCog using a local crowd, the system's accuracy can be kept at a high level in perpetuity, which has been one of the primary challenges in practice for similar prior systems.

2 RELATED WORK

LuzDeploy is related to prior work on (i) navigation in the real world by people with visual impairments, and (ii) mobile crowdsourcing and crowdsensing.

2.1 Real-world accessibility

People with visual impairments learn to navigate independently through extensive Orientation & Mobility (O&M) training. Training on specific routes (*e.g.*, home to work) eases day-to-day navigation, while general purpose navigation strategies are learned for unfamiliar environments [55]. However navigation in unfamiliar areas is complex and presents a barrier to independent travel [19], as understanding the route layout is more difficult without visual access to spatial information. Whereas sighted people visually scan their surroundings for spatial cues like building layout, doorways, or signs, blind pedestrians mostly rely on haptic or audio cues. These cues are typically slow to acquire and typically provide insufficient high-level route and landmark information to navigate in new places. Providing this information en-route make unfamiliar environments much easier to navigate [19, 53]. LuzDeploy provides a method to install and maintain the infrastructure necessary for applications like NavCog [1] to provide this information.

The white cane is the traditional navigation aid for individuals with visual impairments. It extends a person's haptic sensing range, increases their safety during navigation, and acts as an indicator of their visual impairment. However, a white cane cannot provide high-level route guidance. As a result, researchers have studied other approaches to supplement such tools. Outdoors, GPS localization can help support high-level route information if there is adequate coverage [2, 28, 39]. Indoors, however, GPS signals are often too weak; even outdoors, GPS can be limited to a localization accuracy of tens of meters [41]. Navigation assistance techniques that do not rely on GPS often involve adding after-market infrastructure to the environment. A widespread physical add-on is tactile paving [26], which is installed on the floor of the environment to signal routes that people with visual impairments can follow with a white cane. This system works well in persistent installations with a small number of routes. It works less well for complex, temporary, or changing venues, as it is onerous to place or modify tactile paving infrastructure.

Other approaches for providing navigation assistance use sensor-transmitter couplings [16], with an array of transmitters installed in the environment at known locations, and a sensing device carried by the user [4, 35, 46]. In particular, approaches that use Bluetooth low-energy (BLE) beacons [1, 15, 18] are accurate, robust, and have been used to guide people with visual impairments. Most importantly, environments can be augmented with BLE beacon networks after the construction of the physical venue, without requiring structure reconfiguration. However, the installation and calibration procedures for these systems require expertise. Moreover, recruiting and training expert personnel to perform the installation takes time and investment. In particular, for large environments and time critical installation, there is a significant barrier to enroll a sufficient number of skilled installers.

Alternative approaches might seem more feasible, but they also have high set up or maintenance costs. Installing passive RFID tags around the building removes the need for battery replacement. However, the installation needs to be more dense, the user needs to carry higher-power sensing devices, and the tags can still be removed or damaged. Localization based on WiFi does not need additional add-ons beyond WiFi access points [42], and it uses many samples of signal strength to construct a localization model, like Bluetooth approaches. However, the localization is typically not as accurate because there are fewer signals to observe. Computer vision methods utilizing Simultaneous Localization and Mapping (SLAM) do not have any physical infrastructure to install and maintain, but an expert must construct a 3D model of the building's interior and update it whenever something changes. Regardless of the method used, indoor localization techniques suffer from the same basic problem: it

is difficult to install and maintain the system over time. LuzDeploy allows non-experts to do this work, and coordinates their efforts over time.

Minimizing the amount of deployed infrastructure (*e.g.*, beacons) while still providing the required level of localization accuracy is one way to reduce the costs related to the installation and maintenance of a navigation system [3]. An orthogonal approach is to reduce the installation costs by using non-experts to perform the deployment. *Pebble* explored how inexperienced workers can set up navigation infrastructure for the navigation of robots [25]. Results show that, while some individuals could place sensors correctly, others lacked understanding of what they were doing as a whole, leading to confusion in the placement of sensors. LuzDeploy directs crowd members to do specific, simple tasks in the environment, helping to decrease confusion.

In this work, we explore using automated agents to provide a broader picture to non-expert workers about the installation they conduct. Unlike *Pebble*, we explore a real world deployment instead of using the technology within a lab setting. Our work probes coordinating local crowds to participate in the deployment. This approach has the advantage that it could be potentially adopted to bring blind navigation infrastructure into more spaces as any passer-by can potentially join and help in the deployment.

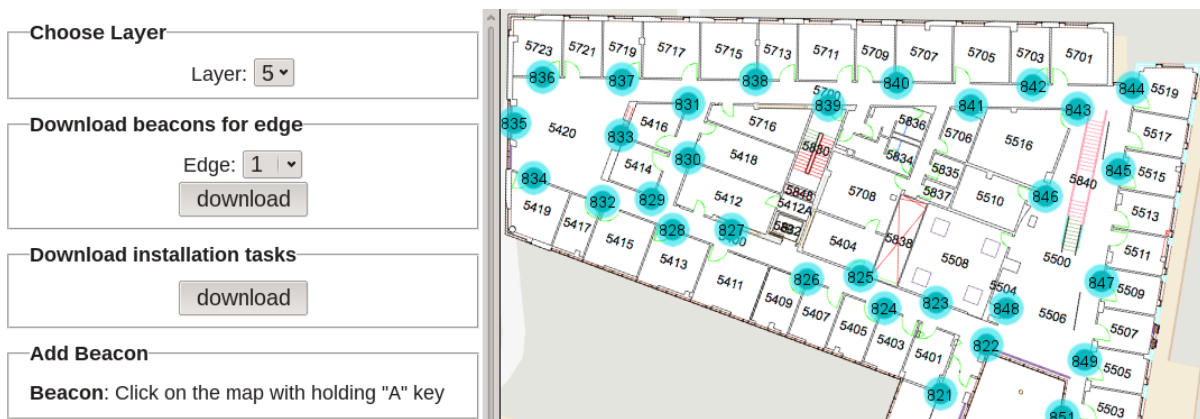


Fig. 2. The LuzDeploy administrator interface allows a new indoor space to be prepared for instrumentation by LuzDeploy. Users select layers (floors) on the left, download lists of beacons for each path segment, or download all beacon installation tasks at once. The right pane allows users to view the floor plan along with the placed beacon markers, or add new markers to the map. This information is then used by the LuzDeploy bot to coordinate the non-expert crowd's effort.

2.2 Mobile Crowdsourcing and Crowdsensing

Physical crowdsourcing involves organizing crowds to perform work in the physical world [33, 45, 51]. TASKer, a recent system for mobile crowdsensing, has shown this approach to be viable on college campuses, distributing tasks to around 900 students [30]. However, this platform focused primarily on data reporting and collection tasks. LuzDeploy combines crowdsensing with physical crowdsourcing to understand both what tasks need to be done, and then coordinate their completion.

Crowdsourced fabrication demonstrated this by guiding participants through step-by-step instructions to manufacture a pavilion [36]. As participants entered the work area, a "foreman engine" guided them through their smart-watch and blinking lights to complete a series of tasks in collaboration with robot arms. The end result was a constructed pavilion made of sticks held in place by string tension. We draw on Lafreniere et al.'s example of step-by-step, just-in-time instructions for the design of LuzDeploy's tasks.

On the topic of volunteer engagement, Heimerl et al. [23] and Goncalves et al. [20] both explored placing physical kiosks throughout universities to motivate workers to complete expert tasks. While these works show how specific local communities could be engaged on the spot to execute work, LuzDeploy uses this approach to recruit participants and then guide them to conduct physical tasks elsewhere in the building.

Crowdwork typically expects workers to just complete the tasks at hand for payment, but micro-volunteering involves volunteers in similar tasks with other motivators like interest in the topic [31]. Because of this, volunteers often know more about the intended goal of a task than in typical crowdwork. Volunteering delivers critical services to communities through collective action, *i.e.*, actions that two or more individuals take to pursue the same collective good. Social computing systems are starting to play an important role in collective action [49], but they are rare because it is difficult to design systems in which a crowd of volunteers can be coordinated to produce large scale change [57]. It is especially difficult to guide crowds to make relevant and useful contributions to the effort [14].

Most volunteering systems focus on making it easy for people to sign-up, input their information, and be matched with relevant tasks [22, 40]. The technological contributions of these systems are primarily matching algorithms that automatically find appropriate tasks for new members based on a measure of their skills and interests [17, 38, 40]. Recently, we have also seen the emergence of research aimed at enhancing volunteer recruitment, by making it easy for people to share volunteer opportunities with friends online [10, 11]. However, systematically guiding volunteers is difficult [24], so many of these systems fail to provide adequate guidance. As a result, many collective efforts fail because substantial work is not produced [24].

Crowdsourcing can work with different kinds of incentives, including monetary, altruism, interest in the subject, *etc.* A number of systems use volunteers for collective action around causes people care about [8]. The “Games with a purpose” paradigm focuses on designing game mechanisms that engage crowds to do useful volunteer work while having fun [13, 54]. These systems mainly focus on guiding volunteers to collectively take action and complete intelligence tasks that are difficult for computers to complete but are relatively easy for humans. Similarly, online micro-volunteering leverages social networks to recruit participants for short tasks [5, 10, 48]. This form of volunteering has recently been used to produce collective action [48], and also solve accessibility problems [7]. For example, Brady et al. [10] leveraged existing social networks to answer questions posed by people with visual impairments about the contents of photos. Most work to date concentrates on using online crowd labor to complete intelligence tasks; LuzDeploy applies micro-volunteering and crowdwork to physical accessibility tasks that would otherwise be difficult to automate.

3 LUZDEPLOY

LuzDeploy is a system that coordinates non-experts via a Facebook bot to use mobile applications designed for non-experts to install and maintain physical localization infrastructure. In the next section, we discuss the design considerations we used in developing LuzDeploy focused on the target use case: the installation and maintenance of NavCog, a navigation assistant for people with visual impairments. We introduce workflows that we developed to guide non-expert volunteers to perform the activities involved in the considered use case. With those in mind, we propose an architecture of the system and describe its components (see Section 3.2).

3.1 Design Considerations

General mobile crowdsourcing tools can take many forms [27, 30], depending on the types of tasks considered and the crowd engaged to complete those tasks. We identified different kinds of tasks involved in the installation, status monitoring, and maintenance activities of NavCog. Based on the identified tasks, we outlined workflows suitable for performing them, as well as methods for engaging a sample target population and promoting participation.

3.1.1 Use Case: Installation and Maintenance of the NavCog System. We envision LuzDeploy as a system for coordinating groups of volunteers to complete mobile, physical crowdsourcing tasks such as installing sensors, collecting signal data, or reporting on infrastructure health. Specifically, we wanted to apply this to the installation and ongoing maintenance of a blind navigation application, NavCog.

A typical NavCog [1] installation, without specialized sampling equipment, currently proceeds as follows:

- (1) An administrator chooses the locations where BLE beacons should be installed (about every 5 meters) and marks them on the building floorplans.
- (2) Each beacon must be installed in the building at the marked locations. The beacons are affixed high on the wall to limit the interference caused by the obstruction of the Bluetooth signals by nearby people and objects. The Minor ID of each placed beacon should match the one intended by the map. The locations of Bluetooth beacons do not have to be exact, and the system will work well even if the beacons are 1-2m off from their designated location.
- (3) Once the beacons are installed, experts use their smartphones to collect many samples of Bluetooth signal data in the building at known locations, also known as “fingerprinting”. Each fingerprint sample measures a few seconds (typically 5) of the signal strength for nearby beacons. When collecting data at fingerprinting locations, experts also mark their positions on a map. This forms the ground truth data for NavCog localization, it is therefore ideal for this estimate to be as accurate as possible.
- (4) The collected fingerprints and the beacons locations are then used to create a model of BLE beacon signals for the building which is used for localization.

After installation, the system must be maintained over time. Besides worrying about the beacon batteries running low, administrators must also look for beacons that have been moved or damaged. The process is the following:

- (1) Every few weeks to months, depending on the overall health of the system (which is hard to estimate precisely), an expert should perform a “sweep” procedure for each hallway of the environment, which consists in sensing the beacons health status with their smartphone. Specifically, this process looks for the beacons expected in that area and marks those that are missing due to damage or battery failure. If a constant flow of NavCog users is present in the chosen environment, this can be done by end users’ navigation applications. However, because NavCog specifically targets infrequent blind users who are unfamiliar with the environment, we do not expect this to happen often. Experts will generally need to complete this process.
- (2) Using the record of missing or damaged beacons collected during the sweeping task, system maintainers then locate and replace or repair the beacons not seen in the sweeping process.

To enable non-expert volunteers to perform the installation and maintenance of NavCog, the identified tasks are segmented into simpler micro-tasks. These include: (i) installing beacons, (ii) collecting fingerprint data, (iii) sweeping a hallway for missing beacons, and (iv) replacing/repairing a broken beacon. The tasks are further defined as sequences of atomic activities which are provided as instructions to the volunteers. Because the first two tasks are completed during installation, a known quantity of work is available all at once. The latter two are maintenance tasks that occur less frequently, although sweeping could be scheduled.

The quality of the completed installation tasks will affect the resulting localization accuracy when users’ navigate with NavCog, but there is some room for error in the setup process. For beacon installation task, a placement error of 1-2 meters would not significantly impair NavCog localization. Similarly, the data sampling tasks will always have some amount of error in the actual location in which the samples were collected versus the reported locations, even when the tasks are performed by experts. Clearly, more error in the reported location of the samples will also lead to a higher localization error in the model. A good measure of quality in this case is to evaluate the localization error from the resulting NavCog model.

3.1.2 Engaging the Crowd. There are many motivations for people to participate in crowd work, especially mobile crowd work [32, 34, 52]. For crowdsourcing or micro-volunteering systems designed to support accessibility, past systems have leveraged money [7, 21], social connections [9], fun games [29], and altruism as motivators [50]. The motivators depend heavily on the installation environment, on the engaged crowd, and on the tasks at hand. For example, in a temporary venue such as a research conference, there are numerous people available for a short time, so the best motivators produce quick and focused bursts of volunteer work. In a university environment, the passers-by tend to be students who are frequently present for a short time between classes or during breaks from work, so they can be engaged for less focused, long-term participation.

With LuzDeploy, we studied the following motivators:

- **Altruism** - Participants were told that the goal of the installation was to make the building accessible to individuals with vision impairments. We believe that this motivator made it easy to recruit new participants as most of the passers-by we engaged with this motivator participated in the installation effort.
- **Interest** - We piqued the interest of potential participants by providing technical information about the installed navigation assistance infrastructure and the NavCog app, as well as about the LuzDeploy system. This motivator works well in a technical context, such as at a university or a scientific conference, in which the passers-by may be knowledgeable and interested in the technical challenges of the approach. We do not anticipate that this motivator would work well in more public contexts, such as a mall.
- **Competitiveness** - This motivator aimed at gamifying the installation process by introducing competitive game mechanics. We achieved this by tracking the results of the participants' work and publishing the participants' scores on a leaderboard in a highly-trafficked area. This aspect was reported to be a strong motivator by one of the most productive participants of Session 2 during an informal interview.
- **Monetary incentives** - We added direct monetary compensation in sessions 3 and 4 to encourage repeat engagement. Given the expense of engaging experts to install and maintain this infrastructure, even fairly generous compensation may be significantly less expensive.

3.2 LuzDeploy Implementation

We targeted our deployment of LuzDeploy to students on a university campus, so we wanted an interface that would make it easy to immediately sign up and receive step-by-step instructions. We developed LuzDeploy as a chat bot using Facebook Messenger¹, as most students already had it installed on their phones. In other participant pools, such as a dedicated weekly volunteer group, drop-in participation might not be an important or desirable feature, and another interface could be preferable. This version of LuzDeploy uses three components:

- (1) The LuzDeploy Bot: The bot is a server application that handles incoming messages and conversation state, and also reminds workers of unfinished tasks. This server maintains a record of workers, placed beacons, and collected Bluetooth fingerprints.
- (2) The LuzDeploy Map: The map is a modified version of the interface supplied by the NavCog system [1], a web view of building floorplans that are overlaid on top of Google Maps. Administrators use this interface to define data for tasks. For workers, this map shows the locations for beacon placement, fingerprint locations, and hallways to be swept.
- (3) The LuzDeploy Data Sampler: For data sampling tasks (*i.e.*, fingerprinting and sweeping), we could not use Facebook alone because it did not allow us to collect Bluetooth data. We instead developed a companion app that would open for those tasks from the bot conversation. This app sends the observed Bluetooth data to the Bot and redirects the user back to the conversation when done.

Figure 3 demonstrates how these components inter-operate to guide users in completing tasks.

¹Specifically, we used Facebook Messenger Send API: <https://developers.facebook.com/docs/messenger-platform/send-api-reference>

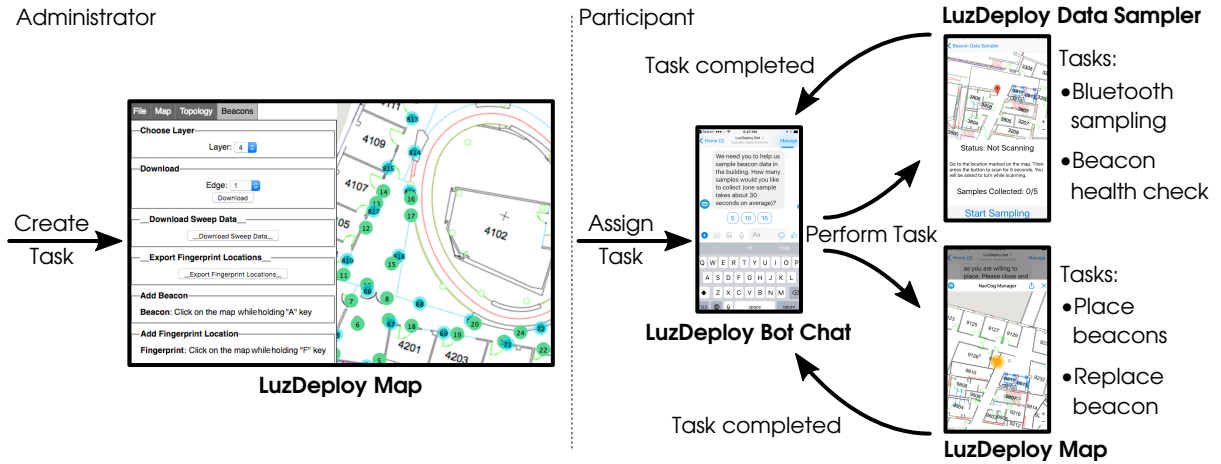


Fig. 3. LuzDeploy components: The administrator uses the online Map interface to create the data for beacon placement, fingerprinting, and sweeping tasks, which are dispatched by the LuzDeploy Bot to participants through Facebook Messenger. Then, the Bot uses either the Map or Data Sampler application to guide the participants to complete the assigned task.

3.3 LuzDeploy Bot

The LuzDeploy Bot is a server that receives incoming messages from workers and responds over the Facebook Messenger platform. The bot manages conversation state with each user, keeping track of current tasks and the state of the deployment.

3.3.1 Task State Machines. The tasks are each implemented as a state machine to control the step-by-step instructions. As commands (see Table 1) or data is sent in the conversation, the bot responds as needed and changes the saved data state.

Command	Action
start	start an assigned task
reject	give up a task if you have one
ask	get another task if you do not have one
earnings	see how much money you have earned
leave	leave the deployment (can rejoin)
help	request a list of commands or personal help

Table 1. List of commands available in LuzDeploy. Instead of a natural language interface, we used simple, one word commands that users can look up.

Each task state can also specify which additional responses it accepts, such as “yes” and “no” for a question, or “number” if that state accepts numerical input. Alternatively, tasks can send buttons to users to perform these responses, or options can be pre-filled for the participants to quickly respond. Data sent from an external application, such as the Data Sampler, can also trigger a state change.

3.3.2 Event mode and Casual mode. Depending on the level of volunteer motivation and recruitment method the administrator may choose to enable one of two different deployment types that affect task distribution: "event-based" or "casual" deployment.

In an "event-based deployment", volunteers are expected to be involved for an extended period of time. This is appropriate if an existing pool of volunteers is being leveraged. When joining the deployment effort, volunteers are immediately given a task to complete. After the completion of each task, another one is immediately assigned and started. This is meant to keep volunteers continuously engaged.

In contrast, "casual deployments" are efforts where volunteers are expected to be involved only sporadically for a few minutes at a time. This is appropriate when there is high volunteer turnover or relatively few tasks to be completed (*e.g.*, ongoing maintenance). In this mode, volunteers ask for a new task when they would like a new one. They receive daily reminders of any unfinished tasks, and are prompted to start a new task if they do not already have one.

In either deployment type, tasks are randomly assigned from what is currently available. For example, if beacons need to be placed on one floor and data needs to be sampled on another, participants have a 50% chance of getting either task type. If users do not want to complete the task they are given, they can "reject" it to abandon the task. Doing this alerts the user that tasks are randomly assigned and they may receive a different one if they ask again. As this is a very simple method of assigning tasks, future versions of LuzDeploy may instead take the performance of volunteers on different task types into account.

3.4 LuzDeploy Map

We modified the open-source NavCog online interface to allow administrators to layout the information used by the bot. Administrators upload pictures of building floorplans, which are then rotated and scaled to fit over their relative locations on Google Maps. They then annotate walkable areas in the building, spots where beacons were to be placed, and locations for fingerprints to be collected.

A read-only version of this map was included with tasks to show workers where beacons were to be placed or samples were to be collected. When the workers clicked a link to the map, it would open with a specific location highlighted for a beacon to be placed, a fingerprint to be collected, or a hallway to be swept. The users were able to move the map around to find the correct location, or zoom in to see a more precise location.

3.5 LuzDeploy Data Sampler

While Facebook Messenger was a good platform choice for deploying beacons, we were unable to collect the Bluetooth signal data needed for our fingerprints without an additional application. The LuzDeploy Data Sampler is an iOS app that can scan for the presence of beacons or collect fingerprint samples. Users were directed to open the app during conversations with the bot, and the clicked link would load all necessary data into the application. Following the map and instructions onscreen, users would press the button to collect data when ready. After the task was complete, the user was directed back to the conversation for additional tasks.

3.6 Tasks

As mentioned previously, each task is implemented as a state machine to control conversation flow over multiple steps and remember its internal state. This is useful for some physical tasks that need special batching logic, such as guiding a user to pick up some beacons, then finding open locations for them that are adjacent. While many bot commands are the same at any moment (see Table 1), this logic that is specific to each task is contained in these state machines.

Below we detail the tasks that LuzDeploy Bot has currently implemented. Example conversations with participants for each task are provided in Appendix A.

3.6.1 Place Beacons. In the Place Beacons task, volunteers are directed to a supply station where they could pick up several beacons. Based on the number of beacons the participant takes, the Bot selects the same number of positions to place them, all close to each other. One by one, the participant is directed to a new location. The user confirms a beacon has not been mistakenly placed at that location already, enters the number of the beacon they are placing, and finally sticks the beacon on the wall. This is repeated until the participant does not have any remaining beacons. If the participant reports that the position already happens to be occupied by a beacon already, the Bot will redirect the participant to a new location.

3.6.2 Bluetooth Sampling. In the Bluetooth Sampling task, the Bot sends the participant a link to the LuzDeploy Data Sampler, pre-loading it with a user-specified number of adjacent fingerprint locations. For each location, the application guides the user using the map and directs them to press the start button. The app then collects Bluetooth RSSI data for five seconds as the participant slowly rotates to create an orientation-invariant sample.

3.6.3 Beacon Health Check. The beacons must be periodically checked to make sure they are still in place and have sufficient battery charge. In this task, participants are given a link to the LuzDeploy Data Sampler which loads a map of a single hallway (*i.e.*, NavCog edge). The participant starts collecting Bluetooth data and walks the length of the hallway, pressing stop when they reach the end. The app reports the presence or absence of beacons that were placed near that hallway.

3.6.4 Replace Beacon. The Replace Beacon task asked volunteers to pick up a new beacon, take it to a location, and place it on the wall in a manner similar to the Place Beacons task. However, if another beacon was already present on the wall, such as if the battery had died, the bot asked the user to return that beacon to the supply station, where administrators could later investigate what failed.

Changing LuzDeploy to support an additional type of multi-step, physical crowdsourcing task just requires creating one Javascript file detailing the states. For example, if we wished to add a task to LuzDeploy to collect pictures around the campus, we may include a step that directs the user to a specific location by guiding them using an embedded map. Once the user indicated they had arrived by issuing a command such as “here”, then the bot would provide instructions for how to take the photo. After the photo had been received by the bot, it would move the task into the generic complete state, allowing the user to take another. Future work should investigate how these tasks could be described without writing additional Javascript code, but still maintain the flexibility of custom logic.

4 DEPLOYMENT ON A UNIVERSITY CAMPUS

We used LuzDeploy to install and begin maintenance of NavCog on the Carnegie Mellon University campus, but this system could be modified for deployment in a variety of public environments. For example: a conference center, hospital, airport, or library.

However, the population in these varying environments are different; students might be in the university every weekday, but the same is not true of travellers in an airport. In some cases, such as a hospital or conference, there may already be a dedicated pool of volunteers that will use the system. On a university campus it may be better to engage students in an installation effort by recruiting them on the quad.

The relationship of a person to the building being instrumented can also have a great effect. We motivated students to participate in our studies partially through altruism, interest in the technical aspects of the project, monetary compensation, and competitiveness with other students. In other venues with different worker pools, these motivators could be completely ineffective. Passengers on a layover in an airport may not be motivated to compete with strangers to a place they will never visit again.

The recruitment of volunteers and motivations for their participation is therefore difficult to proscribe in advance. It is infeasible to design and evaluate a system that would be appropriate for every scenario, so the

manner of recruitment and volunteer motivation is largely left up to the administrator using LuzDeploy. That being said, we did evaluate several variations of features in LuzDeploy that could change based on the volunteer pool being engaged. This evaluation was performed over four different sessions.

4.1 Session 1

In order to prepare the building for the installation, BLE beacons were numbered and their intended locations were marked on the LuzDeploy Map. These locations were manually chosen, based on building floorplans, to yield high localization accuracy for navigation assistance with NavCog. During the first session only, an early LuzDeploy prototype was used that expected each beacon to be placed in these pre-determined locations.

The first session took place on a single day from 4PM to 10PM in the computer science building of the university. Multiple members of the research team recruited students walking through a heavily trafficked hallway near a building entrance. Most workers were able to quickly join the deployment and take a beacon to place in the building. After messaging the bot and filling out the consent form, workers were immediately given a task to install a beacon. This session was meant to be a single-day event, so after each task was completed, a new task was sent to the worker. Event-based deployment continued until the worker stopped participating. At the end of the day, workers were asked to fill out an online survey about their experience with LuzDeploy.

4.1.1 Results. In total, LuzDeploy attracted 89 participants from passerby during this session, with activity peaking around 4PM and 7PM (see Figure 4), which were times of high foot traffic in the building. We noticed that 13 participants did perform the assigned task but did not remember to start the task in the bot chat. These participants and the corresponding 13 beacon placements are therefore not assigned to any of the listed time ranges.

During this session, a total of 99 beacons were placed. 70 participants, which is 78% of those who signed the consent form, placed at least one beacon, while 11 participants placed more than one beacon (see Figure 5a). Because each task required workers to place a specific beacon in a specific spot, volunteers returned to the supply station each time to get the next requested beacon. As a result, each task incurred a large amount of walking overhead. This one-to-one mapping of beacon to placement location also prevented picking up more than one beacon at a time, so most volunteers did not place more than one beacon.

There were 6 participants who signed the consent form but never joined the deployment to get a task, and 13 participants who joined but did not complete a task. Some of these participants encountered technical issues with LuzDeploy Bot or Facebook Messenger, which is why they were unable to continue.

4.1.2 Post-Session Survey. The Bot distributed a survey after the deployment to collect workers' feedback for improvements before the next session. The survey was completed by 21 participants, 23% of those who participated. 85% of these respondents reported prior and frequent previous volunteering activities. The most common reason that survey respondents gave for participating in the study was that they thought the project was interesting or fun (8 of 21, corresponding to 38%). Other reasons are reported in Table 2.

Two-thirds of respondents (14 out of 21) found the tasks short and easy. 16 out of 21 respondents (76%) gave input on how LuzDeploy could be improved. The most common improvement for LuzDeploy, suggested by 9 participants (56% of suggestions), was some method of 'batching' to place multiple beacons at once.




Fig. 4. Number of participants and installed beacons per hour throughout the deployment.




Fig. 5. Number of beacons placed by each participant in each session. Batching was introduced in the second session, so tasks are beacons are segmented by task. While there were fewer participants in the second session, almost the same number of beacons were placed.

Reasons to participate	# of participants	%
Project seemed interesting	8	38%
Participant had free time	4	19%
Invited by a colleague	4	19%
Participant was close by	3	14%
Our stand had food and snacks	3	14%
Interested in accessibility	2	10%

Table 2. Reasons to participate listed by volunteers.

The survey also asked participants to answer why they stopped taking new tasks. Most reported that they had to get back to work or it was late (12 participants, 57%), but some felt they had already contributed enough to the deployment effort and were not motivated to do more (5 participants, 24%). One survey respondent was confused by LuzDeploy Bot’s instructions and could not find help.

In order to understand if LuzDeploy could motivate these workers to help with ongoing maintenance of the beacons, we asked if participants would be willing to check on beacons periodically. 11 (52%) wanted to check on beacons provided they had placed them initially, but only 7 (33%) were willing to spend time inspecting beacons placed by others.

4.1.3 Visual Beacon Inspection. One week after the deployment, a visual inspection of each installed beacon was performed by two members of the research team. NavCog does not need the beacons to be placed exactly at a fixed point [3], so the beacons were inspected to see if they were within 1 meter of the location marked on the building floor plan. Of the 99 beacons deployed in the building, 73 were on the designated map marker or close enough. During inspection, we found that 6 beacons out of those given to the workers were missing and 17 beacons were near the intended location but too far away. Out of these 17, one beacon was found broken apart after the deployment.

4.2 Session 2

To simulate the less frequent and less planned nature of maintenance tasks, our second session occurred over four months.

Based on feedback from the first session, three key improvements were made to beacon placement tasks before the second session. The first was that beacons were no longer tied to a specific location. Instead workers would pick up any beacon, be given a map location, and later tell the bot which one they had placed. This removed the initial time it took to find the correctly numbered beacon. The second improvement reduced overhead time walking to and from the beacon supply point through batching. Workers would take as many beacons as they had time to place, and then the Bot would pick the same number of locations near each other. This minimized the amount of walking required between placing beacons. Finally, tasks were immediately started on assignment for more accurate data collection, although task completion still had to be manually indicated.

Rather than a single-day event, the second session took place over four months, and without a constant presence from the research team. Instead, we recruited participants primarily through fliers, on-line posts on social media, and emails to campus mailing lists. Participants from the first session were invited to complete tasks in the second session, although only one participant placed a beacon in both sessions.

A beacon supply station (Figure 6) was set up in the same hallway as the first session. The stand included a stock of beacons in a locked safe and a tablet. The tablet continuously displayed information on how to join as a volunteer, the overall progress of the deployment, and a leader-board of participants. Each task earned points on completion, and tasks that were batched earned more points.

To simulate the more sporadic nature of maintenance tasks, we utilized a more "casual" task assignment strategy, as volunteers were not given a new task to complete as soon as they finished one. However, workers were either reminded every day about tasks they had not yet completed or about a new task if they didn't have one.

4.2.1 Results. 25 participants placed a total of 97 beacons in the second session, resulting in an average of 3.9 beacons per person (2.7 per task over 36 tasks).

The majority of tasks were completed in the afternoon, although we did see some volunteers in the morning or late at night because the building is open to students at all hours (see Fig 7a). Task reminders were sent out at 12:30PM every day based on a note in prior work [30], but activity did not spike from 12PM - 2PM compared to later in the afternoon.

No beacons were placed on a Saturday, but some were on a Sunday, surprisingly. Overall most of the tasks and beacons are concentrated on Tuesday - Thursday, perhaps due to when classes are scheduled.

4.2.2 Visual Beacon Inspection. After the second session, we again visually inspected the placed beacons to see what the placement accuracy was. We found that 83 beacons had been placed correctly, and



Fig. 6. In the second session, a small end table replaced the presence of the research team. On the table is a lock-box with beacons and a tablet displaying a leaderboard, progress, and sign up information.

Beacon Status	Session 1	Session 2
Placed Correctly	73	83
Off Marker	17	10
Broken	1	2
Not Found	6	2
Total	99	97

Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, Vol. 1, No. 0, Article 99, Publication date: March 2018.

Table 3. The number of beacons correctly placed (or otherwise) from each session. We can see that the majority were placed correctly. Cross-participant verification tasks could catch these cases and move the beacons to the correct location.




Fig. 7. Number of tasks completed and installed beacons by day of week and time of day throughout the second session. The times and days likely mirror student activity in the building, such as classes. Some outliers (e.g., 1AM and Sundays) remain.

only 12 were too far from the intended location or not found. Two beacons placed in this session were found broken during the inspection.

The percentage of beacons placed correctly increased in the second session, and two possible explanations exist for this. A participant who decided to place a batch of beacons when they had the time, instead of when they met the researchers in the hallway, may have taken more care when completing the tasks. Alternatively, because a smaller number of participants placed more beacons in this session, they may have learned how to place them more effectively, gaining expertise over repeated tasks.

In the visual inspections of beacons after each session we noticed that many participants struggled to find the exact location to place the beacons, leading to some beacons being as much as 1-2m off from the specified location. This usually happened in areas with few features, such as bare hallways, where it was not easy to identify the correct beacon location. The participants also tended to “lock on” to salient locations for beacon placement, such as corners or boundaries on walls. The placement of these beacons likely influenced later volunteers, as many beacons in the same area were placed in a similar way.

We observed some other common reasons that resulted in wrong positioning of the beacons:

- **Material matters:** When the beacon was to be placed on glass windows or bulletin boards, many participants instead placed the beacons on a nearby wall.
- **Maps are ambiguous:** Beacons that were to be placed in the middle of a bare wall were often off the map marker. This was especially true in hallways where the only distinguishing map features were interior room boundaries not visible to the volunteer.
- **Maps are not always correct:** Some locations on the map could have been landmarks for participants, but there was not a corresponding location in the building, as the map was out of date.
- **Height was not consistent:** We requested that participants try to place the beacons as high as possible, but some were placed too low. In some areas, participants placed the beacons level with those previously placed, but we observed many instances where they were inconsistent (see Figure 8). The overhead view of the map did not provide guidance on height.




Fig. 8. Beacons placed in poor locations: (a) too low on the wall, (b) on the underside of a light fixture, (c) on the floor.

- **Orientation is not easy to discern:** Beacons on pillars and wall corners were often not facing the direction intended by the map (which did not show an arrow to indicate orientation), as participants just tried to get them near to the landmark.

Less ambiguous map locations will likely lead to deployments with more accurate beacon placements. This adds a second constraint on the administrators when designating the beacon locations for the deployment. Not only must we consider the beacon locations that lead to high accuracy for NavCog, but also which locations are more suitable from the workers' perspective. However, since the highest localization accuracy is actually desired in the proximity of salient environment features, such as turns and doors [3], placing beacons near those elements is beneficial, and ultimately does not impose additional limitations to the task organization. Based on this intuition, in the following sessions, the data collection tasks were designed to correspond to easily identifiable environment features such as doors, pillars, and corners.

4.3 Session 3

Over a two week period, we tested the Bluetooth signal fingerprinting and sweeping tasks. These tasks were deployed in the same manner as the second session, although only volunteers with iOS devices were allowed to complete them. The Facebook bot sent participants a link that opened the LuzDeploy Data Sampler app.

Because our participant pool was constrained to iOS devices and those who would download the companion application, we offered compensation to encourage more continuous participation. Participants were paid \$0.125 per fingerprint sample collected and \$0.40 per hallway swept, which came to an equivalent average hourly rate of \$15.

4.3.1 Results. For fingerprinting, 9 participants collected 1-4 samples at 112 locations, resulting in 185 total samples collected (mean of 20.5 per participant). The samples were collected on 2 different floors, and were completed in batches ranging from 2 to 18 samples (mean of 11.6 per task).

The number of fingerprint samples we collected was impressive for the small number of participants we recruited in such a short period of time. Recruiting a larger group of participants would enable us to quickly perform installations at scale. Future work will investigate how we may quickly recruit more participants and encourage more engagement with LuzDeploy.

Finally, we also wanted to assess how LuzDeploy can use physical crowdsourcing to monitor the system health, in this case by sweeping for the beacons' presence in building hallways. Sweeping is needed only sporadically and there are far fewer hallways than fingerprint locations, so only 9 of these tasks were completed by 3 separate participants (who also collected fingerprint samples in this session). These participants swept a total of 25 beacons, finding all of them reporting correctly.

4.4 Session 4

The third session gave us insight on whether participants could complete data sampling and health check tasks, but we wanted to assess the quality of the data they collected. To do this, we gathered enough data to build a NavCog model over a month-long period. The fourth session was limited to a single floor in the building that was easily accessible to all participants. Eight participants were recruited through a LuzDeploy Bot message and fliers posted in the building, and two were repeat participants. The participants were paid the same amounts as in Session 3.

4.4.1 Results. The eight participants collected a total of 1638 samples across 116 points on a single floor in the building. Almost all of the samples can be attributed to two participants, who collected 1116 and 352 samples, respectively.

To estimate the quality of the collected samples, we compared a NavCog model built with the data from Session 4 to a NavCog model built by experts. The Session 4 model contained 1638 samples at 116 locations on a single floor. The expert-built model was created using the latest NavCog sampling method, which uses a LiDAR to ensure the Bluetooth fingerprints are accurately labelled [47].

Once we constructed the two localization models, we measured the localization accuracy achieved on each of them using a third set of 590 test fingerprints with known ground-truth locations. We then compared the localization results obtained by the two models. In all cases, the localization was performed using single fingerprint samples, instead of the user's trajectory over time. While NavCog makes use of the smartphone's inertial motion unit (IMU) to estimate the user's movement, for this evaluation we did not use movement data. This constraint makes it easier to compare the localization estimate against known ground-truth samples, but it also inflates the error estimates in both models compared to what a user would actually experience when using NavCog.

Figure 9 reports the results of the accuracy evaluation. The model constructed with LuzDeploy data from Session 4 has a higher localization error (median=1.6m, avg=2.5m, std=2.8m) than the expert-constructed model using LiDAR-labelled data (median=1.2m, avg=1.6m, std=1.1m), which is to be expected. However, there is only a 0.4m increase in the median error, meaning the model would still be usable for localization with slightly reduced accuracy. Furthermore, because beacon placement and fingerprint collection points are selected in proximity of turning points, doors, and other features of interest, the decrease in localization accuracy is not expected to have as much impact on the key areas in which accurate localization is required.

5 DISCUSSION

5.1 Streamlining Installation and Worker Participation

The installation process for navigation systems like NavCog is a complex activity, and LuzDeploy serves as a valuable tool to break down the workflow into small tasks. The segmented tasks are easy for non-experts to start and quickly complete, thus making the entire process more accessible. We were able to recruit 127 non-expert individuals who completed 455 tasks, reducing the workload of each single task to several minutes at a time.




Fig. 9. Error average, median, and standard deviation for the expert-built and Session 4 NavCog models.

On the whole, we found that volunteers would quickly drop-in and drop-out of our deployments. Driven by curiosity and the low effort required to start using LuzDeploy, many passerby would participate in the activity, but most volunteers would only complete a single task. One aspect that we explored to improve the efficiency of the volunteer work is to make it possible for the participants to define the amount of activity they are willing to commit to the project. This approach was used after session 1, as participants could complete tasks as their schedule allowed and set how much work each task included. This resulted in a higher task per person ratio, and therefore a more efficient use of our participant pool.

A particularly hard aspect of orchestrating physical crowdsourcing activities was maintaining this constant pool of volunteers. LuzDeploy not only needs to raise immediate interest but also encourage participants to commit to follow-through on further tasks over time. Based on survey feedback in the first session, we believe participants might have felt some investment and responsibility for beacons they personally placed and were therefore more likely to participate in the upkeep of that portion of the infrastructure. In future work, we hope to investigate how personal involvement in the placement of beacons influences the degree of participant retention during maintenance activities.

Additionally, involving workers in the organizational and decision making process could lead to heightened engagement. At the moment, LuzDeploy is designed in a manner similar to existing crowdwork: small, simple, and proscribed tasks. However, prior research suggests that giving workers further autonomy and ability to demonstrate competency would be a good motivator [34]. In many ways, collective action projects enable this, as volunteers and activists organize themselves and set the direction of their movement, sometimes with technological scaffolding [49, 57].

5.2 Worker Incentives

Through our various sessions, we employed four different incentive mechanisms to recruit and continuously engage workers. Altruism in participating in a project to help people with vision impairments and interest in the technical aspects of our projects helped us recruit volunteers in Session 1, as many people were curious about how LuzDeploy and NavCog worked to make the building more accessible. However, we found that this curiosity and initial interest did not translate as well to recruiting participants in the casual deployments of sessions 2. Therefore, we tried to add a competitive aspect using a leaderboard with scores for each participant. While some said this motivated them to complete tasks, a full-scale game built around LuzDeploy could provide better results than this simple aspect. In sessions 3 and 4 we added monetary incentives to both attract participants who fit our criteria and keep them engaged over the long term, which we found to be the case for some of our participants.

As previously mentioned in Section 4, the design of LuzDeploy incentives may not directly translate to other venues or participant pools. In a volunteer pool for a local hospital, perhaps the altruistic goal of making the building more accessible is enough to motivate volunteers. Monetary compensation, or something of similar value, is likely a more general motivator than interest in the technical underpinnings of NavCog. In airports, an airline could offer free rewards through their smartphone application for people who completed tasks in the terminal. However, airline passengers would not have the same connection to the building as students. Spatial augmented-reality games, like Pokemon GO² or Ingress³, have also shown that people can be mobilized to specific locations and engaged over time [12], so this approach may also make sense for the general public. In future sessions, it would be beneficial to collect data with these different worker pools and venues to see how the incentives and design of LuzDeploy perform across different populations and venues.

²<https://www.pokemongo.com/>

³<https://www.ingress.com/>

5.3 Beacon Placement

In physical crowdsourcing tasks that first require the workers to reach the location in which the task is performed, ambiguity in guiding the participants to the task location may result in difficulties in performing the task, and therefore in lower work quality. While NavCog can tolerate a certain degree of imprecision in the installation of beacons and data collection, these directly impact the localization accuracy of the resulting NavCog infrastructure. In session 4, defining crowdsourcing tasks in relation to visually salient environment features was crucial to enable correct guidance during task dispatching, and thus guarantee a higher quality of the resulting infrastructure.

Another insight derived from the sessions that involved beacon placement is that people are generally not accustomed to others changing physical infrastructure. In some cases attaching a physical object to a wall in a building that is not one's own space might even be considered vandalism. Clearly, issues of trust and authority will need to be considered in other deployments and the development of future systems. This brings up questions for future work: Should rogue deployments of beacons in places where permission has not been granted be explicitly discouraged? How will workers know the difference?

5.4 Event and Casual Deployments

Because natural battery failure does not usually occur in the time span we ran the study over (7 months), we approximated this aspect in our second, third, and fourth sessions using a "casual" deployment type. In the second session, this resulted in fewer participants involved in the deployment, but those who were involved placed more beacons overall and more beacons correctly. We hypothesized that this might be due to the participants gaining experience over repeated beacon placements, or because they took more care when they completed tasks on their own schedules.

The benefit of our "event-based" deployment in the first session was involvement from many participants, leading us to believe that using this for the initial installation could greatly reduce the individual workloads of beacon placement and fingerprinting. Afterwards, a "casual" mode of task distribution could be used for ongoing sweeping and beacon replacement.

5.5 Quality of Completed Tasks

The quality of tasks completed by the participants during installation is important to monitor, as this will affect the localization experience by end users of the NavCog system. Overall, during the beacon placement tasks, we found that the majority of beacons were placed correctly, with between 10-20% being placed too far from the marked location. This misplacement is undesirable, but has a minimal effect on NavCog's localization accuracy, as the fingerprint samples are more important for the localization system than the absolute position of the beacons. We found that a NavCog model generated with the samples collected by our crowd workers was not as accurate as an expert model, but still provided usable localization.

Quality control mechanisms could increase the accuracy of the NavCog model generated with crowdsourced samples. Verification tasks would ask workers to manually check each other's results to ensure beacon placement is as accurate as possible. Automatic methods, such as outlier detection, could identify workers who submit invalid data or low-quality fingerprint samples. By combining these two methods, LuzDeploy Bot could direct workers to collect more quality data in areas with high localization error rates. After further work to ensure quality, we may see LuzDeploy data get closer to expert-installation accuracy levels.

We must consider whether the installation completed by the crowd workers was worth time and cost expended, compared to an expert installation. Over the course of our third and fourth sessions, we paid our participants nearly \$300, which is likely less than a team of experts would charge. They did not require any training beyond the instructions that LuzDeploy Bot delivered, and were already available in the building, instead of having experts travel to our location. However, the crowd work was spread out over weeks, and experts would likely

focus on completing the installation in much less time. Additionally, while the system created by crowdworkers was sufficiently accurate to enable navigation assistance with NavCog, the expert model achieved a higher localization accuracy than the one using data collected by the crowd workers. If an organization would like quick installation and the highest accuracy, it may make sense to spend more money and hire an expert to come in. However, if the goal is to install indoor navigation systems as widely as possible, crowdwork may be a viable option.

5.6 Future Work

Beyond the areas of further investigation already outlined in previous sections, there are several improvements for the LuzDeploy system that could impact future deployments. Our participant pool is currently limited to workers with an iOS device, but by porting the LuzDeploy Data Sampler to the Android operating system, we could engage android users in all tasks, including data sampling. Additionally, the LuzDeploy Map and task descriptions could be improved to provide more accurate guidance, thus reducing user error in task execution. Finally, implementing the quality control and verification mechanisms, mentioned previously, would provide participants with needed feedback so that they may improve their work.

Our deployment of LuzDeploy was intentionally not a controlled laboratory experiment, as we wanted to collect data in a large-scale, realistic setting. However, data from a controlled study could be valuable, so we would like to conduct experiments in small scenarios where more controlled data analysis could be performed. Specifically, controlled variables such as the timing of task reminders, instruction text, and deployment type (*i.e.*, event vs. casual) could be compared for statistically significant differences in outcomes. A testbed for these controlled experiments could also provide an overall comparison of physical crowdsourcing with other installation methods, such as experts or robotic installations.

More generally, we believe that LuzDeploy advocates future work in understanding and supporting adoption of assistive technology longer into its deployment than has been done previously. We know that many tools are abandoned because they are too difficult to learn, break and the user does not know how to fix them, or gradually degrade in effectiveness [44]. Future work may explore how crowdsourcing and collective action can make technology feasible that was previously too difficult to deploy because it requires on-going maintenance.

6 CONCLUSION

The technology to create accurate and usable indoor navigation aids for people with visual impairments is here. However, indoor navigation solutions that require extensive installation and maintenance will not become widespread unless the process is simplified. In this paper, we have introduced LuzDeploy, a system that coordinates the work of non-experts in installing and maintaining needed infrastructure via a Facebook Bot using supporting mobile applications. On-site workers, such as the students in our studies, complete small tasks in support of this broader goal without needing special expertise. We demonstrated two different methods of using LuzDeploy: a casual and event approach for task distribution, which are appropriate for different venues and participant populations. Once the infrastructure for these navigation systems are installed, LuzDeploy ensures the system is kept operational by preventing decay due to broken or powerless beacons. In our study, LuzDeploy successfully coordinated 127 participants to install indoor localization infrastructure in a 7-story building and maintain it over several months. As long-term maintenance is a major limiting factor in real-world adoption of current technology, LuzDeploy is a valuable aid for moving this technology out of the lab setting. By overcoming the logistical limitations of these installations, we can finally create widespread and usable indoor navigation systems for people with visual impairments.

ACKNOWLEDGMENTS

We would like to thank all of the wonderful participants in this study, as it could not have been completed without their enthusiasm and hard work.

REFERENCES

- [1] Dragan Ahmetovic, Cole Gleason, Chengxiong Ruan, Kris Kitani, Hironobu Takagi, and Chieko Asakawa. 2016. NavCog: A Navigational Cognitive Assistant for the Blind. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '16)*. ACM.
- [2] Dragan Ahmetovic, Roberto Manduchi, James M Coughlan, and Sergio Mascetti. 2015. Zebra Crossing Spotter: Automatic Population of Spatial Databases for Increased Safety of Blind Travelers. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*. ACM, 251–258.
- [3] Dragan Ahmetovic, Masayuki Murata, Cole Gleason, Erin Brady, Hironobu Takagi, Kris Kitani, and Chieko Asakawa. 2017. Achieving Practical and Accurate Indoor Navigation for People with Visual Impairments. In *Web for All Conference*. ACM.
- [4] Paramvir Bahl and Venkata N Padmanabhan. 2000. RADAR: An in-building RF-based user location and tracking system. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*. Ieee.
- [5] Michael Bernstein, Mike Bright, Ed Cutrell, Steven Dow, Elizabeth Gerber, Anupam Jain, and Anand Kulkarni. 2013. Micro-volunteering: helping the helpers in development. In *Proceedings of the 2013 conference on Computer supported cooperative work companion*. ACM.
- [6] Jeffrey P. Bigham, Michael Bernstein, and Eytan Adar. 2015. Human-Computer Interaction and Collective Intelligence. (2015). <https://docs.google.com/file/d/0B4-bDrtyS3lXdFJIRXk2bGdpQzg/edit>
- [7] Jeffrey P Bigham, Chandrika Jayant, Hanjie Ji, Greg Little, Andrew Miller, Robert C Miller, Robin Miller, Aubrey Tatarowicz, Brandyn White, Samuel White, and others. 2010. VizWiz: nearly real-time answers to visual questions. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*. ACM.
- [8] KD Borne and Zooniverse Team. 2011. The Zooniverse: A framework for knowledge discovery from citizen science data. In *AGU Fall Meeting Abstracts*.
- [9] E Brady. 2015. Using social microvolunteering to answer visual questions from blind users. *ACM SIGACCESS Accessibility and Computing* (2015).
- [10] Erin Brady, Meredith Ringel Morris, and Jeffrey P Bigham. Gauging receptiveness to social microvolunteering. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM.
- [11] Justin Cheng and Michael Bernstein. 2014. Catalyst: triggering collective action with thresholds. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. ACM, 1211–1221.
- [12] Ashley Colley, Jacob Thebault-Spieker, Allen Yilun Lin, Donald Degraen, Benjamin Fischman, Jonna Häkkinä, Kate Kuehl, Valentina Nisi, Nuno Jardim Nunes, Nina Wenig, and others. 2017. The geography of Pokémon GO: beneficial and problematic effects on places and movement. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 1179–1192.
- [13] Seth Cooper, Firas Khatib, Adrien Treuille, Janos Barbero, Jeehyung Lee, Michael Beenen, Andrew Leaver-Fay, David Baker, Zoran Popović, and others. 2010. Predicting protein structures with a multiplayer online game. *Nature* (2010).
- [14] Steven Dow, Anand Kulkarni, Scott Klemmer, and Björn Hartmann. 2012. Shepherding the crowd yields better work. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. ACM, 1013–1022.
- [15] Karen Duarte, José Cecilio, Jorge Sá Silva, and Pedro Furtado. 2014. Information and assisted navigation system for blind people. In *Proceedings of the 8th International Conference on Sensing Technology*. 470–473.
- [16] Navid Fallah, Ilias Apostolopoulos, Kostas Bekris, and Eelke Folmer. 2013. Indoor Human Navigation Systems: A Survey. *Interacting with Computers* 25, 1 (2013), 21. DOI: <http://dx.doi.org/10.1093/iwc/iws010>
- [17] Maryam Fazel-Zarandi and Mark S Fox. 2009. Semantic matchmaking for job recruitment: an ontology-based hybrid approach. In *Proceedings of the 8th International Semantic Web Conference*.
- [18] Aura Ganz, Siddhesh Rajan Gandhi, Carole Wilson, and Gary Mullett. 2010. INSIGHT: RFID and Bluetooth enabled automated space for the blind and visually impaired. In *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*. IEEE, IEEE, 331–334.
- [19] Reginald G Golledge, Roberta L Klatzky, and Jack M Loomis. 1996. Cognitive mapping and wayfinding by adults without vision. In *The construction of cognitive maps*. Springer, 215–246.
- [20] Jorge Goncalves, Denzil Ferreira, Simo Hosio, Yong Liu, Jakob Rogstadius, Hannu Kukka, and Vassilis Kostakos. 2013. Crowdsourcing on the Spot: Altruistic Use of Public Displays, Feasibility, Performance, and Behaviours. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM.
- [21] Kotaro Hara, Vicki Le, and Jon Froehlich. 2013. *Combining crowdsourcing and google street view to identify street-level accessibility problems*. ACM, New York, New York, USA.

- [22] Fuad Mire Hassan, Imran Ghani, Muhammad Faheem, and Abdirahman Ali Hajji. 2012. Ontology matching approaches for erecruitment. *International Journal of Computer Applications* 51, 2 (2012).
- [23] Kurtis Heimerl, Brian Gawalt, Kuang Chen, Tapan Parikh, and Björn Hartmann. 2012. CommunitySourcing: Engaging Local Crowds to Perform Expert Work via Physical Kiosks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM.
- [24] Benjamin Mako Hill. 2014. Almost Wikipedia: What eight early online collaborative encyclopedia projects reveal about the mechanisms of collective action. In *Working Paper*. ACM.
- [25] Kentaro Ishii, Haipeng Mi, Lei Ma, Natsuda Laokulrat, Masahiko Inami, and Takeo Igarashi. 2013. Pebbles: User-Configurable Device Network for Robot Navigation. In *14th International Conference on Human-Computer Interaction (INTERACT) (Human-Computer Interaction – INTERACT 2013)*, Paula Kotzé, Gary Marsden, Gitte Lindgaard, Janet Wesson, and Marco Winckler (Eds.), Vol. LNCS-8118. Springer, Cape Town, South Africa, 420–436. DOI: http://dx.doi.org/10.1007/978-3-642-40480-1_28 Part 9: Humans and Robots.
- [26] Hideyuki Iwahashi. 1983. *Toward white wave - Story of Seiichi Miyake (in Japanese)*. Traffic Safety Research Center.
- [27] Nikita Jaiman, Randy Tandriansyah, Thivya Kandappu, and Archan Misra. 2016. *A campus-scale mobile crowd-tasking platform*. ACM, New York, New York, USA.
- [28] Hernisa Kacorri, Sergio Mascetti, Andrea Gerino, Dragan Ahmetovic, Hironobu Takagi, and Chieko Asakawa. 2016. Supporting Orientation of People with Visual Impairment: Analysis of Large Scale Usage Data. In *International ACM SIGACCESS Conference on Computers and Accessibility*. ACM.
- [29] Hernisa Kacorri, Kaoru Shinkawa, and Shin Saito. 2014. *Introducing game elements in crowdsourced video captioning by non-experts*. ACM, New York, New York, USA.
- [30] Thivya Kandappu, Nikita Jaiman, Randy Tandriansyah, Archan Misra, Shih-Fen Cheng, Cen Chen, Hoong Chuin Lau, Deepthi Chander, and Koustuv Dasgupta. 2016. *TASKer: behavioral insights via campus-based experimental mobile crowd-sourcing*. ACM, New York, New York, USA.
- [31] Bob Kanefsky, Nadine G Barlow, and Virginia C Gulick. 2001. Can distributed volunteers accomplish massive data analysis tasks. *Lunar and Planetary Science* 1 (2001).
- [32] N Kaufmann, T Schulze, and D Veit. 2011. More than fun and money. Worker Motivation in Crowdsourcing-A Study on Mechanical Turk. *AMCIS* (2011).
- [33] Yongsung Kim, Emily Harburg, Shana Azria, Elizabeth Gerber, Darren Gergle, and Haoqi Zhang. 2015. Enabling physical crowdsourcing on-the-go with context-sensitive notifications. In *Third AAAI Conference on Human Computation and Crowdsourcing*.
- [34] Masatomo Kobayashi, Shoma Arita, Toshinari Itoko, Shin Saito, and Hironobu Takagi. 2015. Motivating Multi-Generational Crowd Workers in Social-Purpose Work. In *the 18th ACM Conference*. ACM Press, New York, New York, USA, 1813–1824.
- [35] Nisarg Kothari, Balajee Kannan, Evan D Glasgwo, and M Bernardine Dias. 2012. Robust indoor localization on a commercial smart phone. *Procedia Computer Science* (2012).
- [36] Benjamin Lafreniere, Marcelo H Coelho, Nicholas Cote, Steven Li, Andy Nogueira, Long Nguyen, Tobias Schwinn, James Stoddart, David Thomasson, Ray Wang, Thomas White, Tovi Grossman, David Benjamin, Maurice Conti, Achim Menges, George Fitzmaurice, Fraser Anderson, Justin Matejka, Heather Kerrick, Danil Nagy, Lauren Vasey, Evan Atherton, and Nicholas Beirne. 2016. Crowdsourced Fabrication. In *the 29th Annual Symposium*. ACM Press, New York, New York, USA, 15–28.
- [37] Patrick Lazik, Niranjini Rajagopal, Oliver Shih, Bruno Sinopoli, and Anthony Rowe. 2015. *ALPS: A Bluetooth and Ultrasound Platform for Mapping and Localization*. ACM, New York, New York, USA.
- [38] Hexin Lv and Bin Zhu. 2006. Skill ontology-based semantic model and its matching algorithm. In *2006 7th International Conference on Computer-Aided Industrial Design and Conceptual Design*. IEEE, 1–4.
- [39] Roberto Manduchi. 2012. Mobile vision as assistive technology for the blind: An experimental study. In *International Conference on Computers for Handicapped Persons*. Springer, 9–16.
- [40] Malgorzata Mochol, Radoslaw Oldakowski, and Ralf Heese. 2004. Ontology based Recruitment Process.. In *GI Jahrestagung (2)*. 198–202.
- [41] Marko Modsching, Ronny Kramer, and Klaus ten Hagen. 2006. Field trial on GPS Accuracy in a medium size city: The influence of built-up. In *3rd workshop on positioning, navigation and communication*. 209–218.
- [42] J Rajamäki, P Viinikainen, and J Tuomisto. 2007. LaureaPOP indoor navigation service for the visually impaired in a WLAN environment. In *Proceedings of the 6th ...*
- [43] L Ran, S Helal, and S Moore. 2004. Drishti: an integrated indoor/outdoor blind navigation system and service. In *Second IEEE Annual Conference on Pervasive Computing and Communications, 2004. Proceedings of the*. IEEE, 23–30.
- [44] Marti L Riemer-Reiss and Robbyn R Wacker. 2000. Factors associated with assistive technology discontinuance among individuals with disabilities. *Journal of Rehabilitation* 66, 3 (2000), 44.
- [45] Adam Sadilek, John Krumm, and Eric Horvitz. 2013. Crowdphysics: Planned and opportunistic crowdsourcing for physical tasks. *SEA* (2013).
- [46] Shigeru Saito, Atsushi Hiyama, Tomohiro Tanikawa, and Michitaka Hirose. 2007. Indoor marker-based localization using coded seamless pattern for interior decoration. In *2007 IEEE Virtual Reality Conference*. IEEE.

- [47] Daisuke Sato, Uran Oh, Kakuya Naito, Hironobu Takagi, Kris Kitani, and Chieko Asakawa. 2017. NavCog3: An Evaluation of a Smartphone-Based Blind Indoor Navigation Assistant with Semantic Features in a Large-Scale Environment. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '17)*. ACM, New York, NY, USA, 270–279. DOI: <http://dx.doi.org/10.1145/3132525.3132535>
- [48] Saiph Savage, Andres Monroy-Hernandez, and Tobias Hollerer. 2016. Botivist: Calling Volunteers to Action using Online Bots. In *Proceedings of the 2016 conference on Computer supported cooperative work*. ACM, 839–848.
- [49] Aaron Shaw, Haoqi Zhang, Andres Monroy-Hernandez, Sean Munson, Benjamin Mako Hill, Elizabeth Gerber, Peter Kinnaird, and Patrick Minder. 2014. Computer supported collective action. *interactions* 21, 2 (2014), 74–77.
- [50] Hironobu Takagi, Shinya Kawanaka, Masatomo Kobayashi, Daisuke Sato, and Chieko Asakawa. 2009. Collaborative Web Accessibility Improvement: Challenges and Possibilities. (Aug. 2009), 1–8.
- [51] TaskRabbit. 2016. TaskRabbit connects you to safe and reliable help in your neighborhood. Website. (20 September 2016). Retrieved September 20, 2016 from <https://www.taskrabbit.com/>.
- [52] Rannie Teodoro, Pinar Ozturk, Mor Naaman, Winter Mason, and Janne Lindqvist. 2014. *The motivations and experiences of the on-demand mobile workforce*. ACM, New York, New York, USA.
- [53] D W Tuttle and N R Tuttle. 2004. Self-esteem and adjusting with blindness: The process of responding to life's demands.
- [54] Luis Von Ahn. 2006. Games with a purpose. *Computer* (2006).
- [55] William R Wiener, Richard L Welsh, and Bruce B Blasch. 2010. *Foundations of orientation and mobility*. American Foundation for the Blind.
- [56] Scooter Willis and Sumi Helal. 2005. RFID information grid for blind navigation and wayfinding. In *Wearable Computers, 2005. Proceedings. Ninth IEEE International Symposium on*. IEEE, IEEE, 34–37.
- [57] Haoqi Zhang, Andrés Monroy-Hernández, Aaron D Shaw, Sean A Munson, Elizabeth M Gerber, Benjamin Mako Hill, Peter Kinnaird, Shelly Diane Farnham, and Patrick Minder. 2014. WeDo: End-To-End Computer Supported Collective Action.. In *ICWSM*.

A APPENDIX A: EXAMPLE PARTICIPANT CONVERSATIONS

Below are example conversation with participants for their first interaction with LuzDeploy Bot, as well as their interactions for the beacon placement, data sampling, and health check tasks.

A.1 Getting started

LuzDeploy Bot: Hi! [P1], I am the LuzDeploy bot. To continue you must complete the following consent form. [Button: "Open Consent Form"]
LuzDeploy Bot: Which deployment would you like to join? [Button: "Gates Hillman Center"]
LuzDeploy Bot: Great! Welcome to the Gates Hillman Center deployment! Say 'ask' for a new task. [Button: "ask"]

A.2 Beacon Placement

Participant: ask
LuzDeploy Bot: In this task you will place beacons in the environment that will be used by people with visual impairments to navigate. Please go to the Supply Station (Gates 5th floor near the bridge exit). Tell me when you are 'there'. [Button: "there"]
Participant: there
LuzDeploy Bot: Great! To open the lockbox, type the code 020217, then #, then turn the switch. Now grab as many beacons as you are willing to place. Please close and lock the box. Tell me how many you took (you can press a button or type a number). [Buttons: "1", "3", "5"]
Participant: 1
LuzDeploy Bot: You have 1 beacon to place. Please go to the location marked on the map below. [Button: "Open Map"]
LuzDeploy Bot: Tell me when you are 'there'! [Button: "there"]
Participant: there
LuzDeploy Bot: Is there already a beacon placed on the wall there? [Buttons: "no", "yes"]
Participant: no
LuzDeploy Bot: What is the number on the back of one of the beacons you have?
Participant: 682
LuzDeploy Bot: The beacon number is 682, correct? [Buttons: "yes", "no"]
Participant: yes
LuzDeploy Bot: Place the beacon high on the wall (you can double check using the map), and try to make it look neat. Don't put it on signs, door frames, or light fixtures. [Button: "Open Map"]
LuzDeploy Bot: Tell me when you are 'done'! [Button: "done"]
Participant: done
LuzDeploy Bot: Thanks! There are more tasks available! Say 'ask' to get another.

A.3 Fingerprint Collection

Participant: ask
LuzDeploy Bot: We need you to help us sample beacon data in the building. How many samples would you like to collect (one sample takes about 30 seconds on average)?
Participant: 15
LuzDeploy Bot: Please open the LuzDeploy app below and follow the instructions. Let me know when you are 'done'! [Button: "Open LuzDeploy"]

LuzDeploy Map: [The following instructions are presented in the application alongside maps depicting the location of the sample:] Go to the location marked on the map. Then press the button to scan for 5 seconds. You will be asked to rotate while scanning.

Participant: [participant presses button]

LuzDeploy Map: Now turn around in place slowly until the timer runs out. [5 second timer starts]

LuzDeploy Map: [This repeats until all samples have been collected. The participant is then directed back to the conversation with LuzDeploy Bot.]

LuzDeploy Bot: Thanks, you earned \$1.88. There are more tasks available! Say 'ask' to get another.

A.4 Health Check

Participant: ask

LuzDeploy Bot: We need you to help us check which beacons are not working in the building. Please open the LuzDeploy app below and follow the instructions. Let me know when you are 'done'!
[Button: "Open LuzDeploy"]

LuzDeploy Map: [The following instructions are presented in the application alongside maps depicting the location of the sample:] Go to node 183 in the map above. Then press the start button below.

Participant: [participant presses button]

LuzDeploy Map: Now walk to node 185 at the end of the red path. Once there, press the stop button below.

Participant: [participant presses button and is redirected back to the conversation with LuzDeploy Bot]

LuzDeploy Bot: Thanks, you earned \$0.40. There are no more tasks available right now.

Received May 2017; revised November 2017; accepted January 2018