

Integrating ODK Scan into the Community Health Worker Supply Chain in Mozambique

Nicola Dell¹, Jessica Crawford², Nathan Breit¹, Timóteo Chaluco²,
Aida Coelho², Joseph McCord³, and Gaetano Borriello¹

¹University of Washington
{nixdell, nbdate, gaetano}
@cs.washington.edu

²VillageReach
{jessica.crawford,
timoteo.chaluco, aida.coelho}
@villagereach.org

³John Snow, Inc.
USAID | DELIVER PROJECT
jmccord@jsi.com

ABSTRACT

We describe our experiences integrating ODK Scan into the community health worker (CHW) supply chain in Mozambique. ODK Scan is a mobile application that uses computer vision techniques to digitize data from paper forms. The application automatically classifies machine-readable data types, like bubbles and checkboxes, and assists users with the manual entry of handwritten text and numbers. We designed an intervention that uses paper forms in conjunction with ODK Scan to monitor CHW usage of essential health commodities, finding that the application is capable of providing supervisors and stakeholders with important information regarding health commodity availability in the field. Specifically, we (1) detail our experiences integrating ODK Scan into the health worker supply chain in Mozambique, with particular emphasis on the critical (and often under-reported) role of practitioners; (2) evaluate the impact of the technology at multiple levels of the information hierarchy, providing quantitative and qualitative data that exposes the benefits, challenges and limitations of the technology; and (3) share lessons learned and provide actionable guidance to researchers and practitioners interested in ODK Scan or other systems that bridge the gap between paper-based and digital data collection.

Categories and Subject Descriptors

H.5.2 [Information Interfaces and Presentation (e.g., HCI)]: User Interfaces; I.5.4 [Applications]: Computer vision

General Terms

Design, Measurement, Human Factors

Keywords

ICTD; DEV; HCI4D; computing for development; paper forms; OMR; computer vision; cell phone; smartphone, data entry.



Figure 1: A health district supervisor using ODK Scan to digitize health worker supply chain data.

1. INTRODUCTION

The critical lack of trained medical professionals in developing countries has resulted in the establishment of numerous community health programs in an effort to expand access to health services among vulnerable populations [15]. Community health workers (CHWs) are members of the community who typically receive a small amount of training and who then assess, treat and refer patients in their community according to established health protocols. In addition to seeing patients, health workers are typically also responsible for collecting monitoring and evaluation data about the health programs in which they are engaged.

For a treatment-based community health program to be successful, it is essential that CHWs receive a reliable and uninterrupted supply of health commodities, including diagnostic tests, basic medicines and contraceptives, so that they possess the medical supplies required to treat patients effectively [1]. Securing this uninterrupted supply of commodities requires a strong supply chain to ensure that the right quantities of the right products are available at the right time, place and condition, and for the right cost. However, limited communications and transport infrastructure in many developing countries make it difficult to ensure that medical supplies are ordered, received and distributed on time. In addition, as in Mozambique, there is often no standardized system for collecting and reporting of logistics data, particularly regarding the consumption and delivery of CHW commodities. The lack of such a reporting system leads to CHWs running out of supplies and being unable to treat patients, and

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.

ICTD 2013, December 07-10 2013, Cape Town, South Africa.

Copyright 2013 ACM 978-1-4503-1906-5/13/12, \$15.00.

<http://dx.doi.org/10.1145/2516604.2516611>

supervisors and partners at higher levels of the information hierarchy not knowing CHW consumption patterns or stock status.

To address these challenges, we designed a standardized data collection and reporting system to track CHW consumption of health commodities in Mozambique. In our intervention, CHWs collect data regarding the commodities that they dispense using cheap, familiar and easy-to-use paper forms. CHWs bring these paper forms to their supervisors during their monthly supervisory meeting. As shown in Figure 1, the supervisors use a mobile application called ODK Scan [9] to digitize the CHW consumption data and transmit it to a central server where it is available for immediate viewing and analysis by other stakeholders. ODK Scan is an Android application that uses a smartphone camera and computer vision techniques to digitize data from paper forms. ODK Scan is capable of automatically classifying bubble and checkbox data types and assists users with the manual entry of handwritten text and numbers. We chose ODK Scan for our intervention since we wanted to allow the CHWs to use paper forms, but also wanted to quickly digitize, aggregate and disseminate the collected data.

This paper describes our experiences integrating ODK Scan into the CHW supply chain in two districts in Maputo province, Mozambique. We document the impact of the technology at multiple levels of the information hierarchy: the NGO practitioners who designed the ODK Scan compatible forms, and set up and maintained the system; the NGO field staff who trained and supervised the users of the application; the CHWs who recorded their consumption of medical supplies on the paper forms; and the health district supervisors who digitized and reported the data using ODK Scan. In close consultation with local partners, this activity was led by Task Orders 4 and 7 of the USAID | DELIVER PROJECT [3] in collaboration with the non-profit VillageReach [22].

In describing the integration of ODK Scan into the CHW supply chain, we pay particular attention to the critical and often under-reported role of practitioners in deploying a new technology in the field. Although many ICTD papers present new systems that have been successfully tested with end users, few of these papers consider what it takes for practitioners, typically NGOs or other organizations, to take a new technology, modify it to fit their use-case, and deploy, support and sustain it in a low-resource environment. This lack of attention on practitioners often leads to a mismatch in expectations, resulting in practitioners who want to use new technologies but who do not understand the technical expertise or amount of effort required to successfully integrate the technology into an information ecosystem. We hope that our approach will inspire more work that focuses on strengthening the link between researchers and practitioners in ICTD.

This paper makes three primary contributions:

- (1) We detail our experiences integrating ODK Scan into the health worker supply chain in Mozambique, with particular emphasis on the critical role of practitioners;
- (2) We evaluate the impact of the technology at multiple levels of the information hierarchy, providing quantitative and qualitative data that exposes the benefits, challenges and limitations of the technology; and
- (3) We share lessons learned and provide actionable guidance to researchers and practitioners interested in deploying ODK Scan or other systems that bridge the gap between paper-based and digital data collection.

2. RELATED WORK

The ubiquitous use of paper forms for data collection in the developing world has resulted in a large amount of research that focuses on extracting digital data from paper. CAM [16] was one of the first systems to contribute to this design space. CAM uses visual codes and a camera-phone to assist users with data entry from paper forms. CAM is a powerful tool that can handle a variety of data types, although there are a number of features that distinguish CAM from our work. First, although the visual codes help to identify each field correctly, users are still required to type all of the data into the phone. In contrast, ODK Scan is able to automatically interpret bubbles and checkboxes so that users do not have to enter these data types manually. In addition, the evaluation of CAM [16] focused on documenting the experiences of end-users, whereas we document user experiences at all levels of the information hierarchy.

Shreddr [7] or Captricity [5] has also contributed substantially to this design space. Images of forms are segmented into form fields that are entered by people via crowd-sourcing. Although Shreddr can handle a wide variety of data types, the system does not yet leverage the direct machine readability of certain data types, like bubbles and checkboxes, although it does have a method of doing machine learning based on a sample of the crowd-sourced answers. Shreddr also does not provide organizations with the option of distributing the data entry to its own workers instead of crowd-sourced workers, which may be preferable for a variety of political and security reasons. Finally, a reliable Internet connection and sufficient bandwidth are required for the effective use of a crowd-sourcing platform, which is problematic for many organizations in developing countries.

Local Ground [21] allows users to annotate paper maps using pens and markers. The maps are then scanned, and user markings overlaid on existing online maps to aid planning decisions. Unlike ODK Scan, Local Ground treats the user markings on the paper as an image layer and is currently not capable of reading or making sense of them. Ratan et al. [18] present a financial record management system built on a low-cost digital slate device. Testing of the paper-pen slate system showed that data can be collected more quickly with fewer incorrect entries, and users also liked having paper evidence of their transactions. Unfortunately, the purchase and maintenance of specialized slate devices hindered the scalability and sustainability of the system [14]. In addition, the study focuses primarily on the experiences of the slate end-users. The PartoPen [20] aims to improve maternal health outcomes through careful labor monitoring. The system uses a digital pen that is capable of triggering audio alerts to remind health care workers to take routine patient measurements at specified time intervals. The system can also interpret marks on the paper and alert the attendant of potential complications. However, extracting data from the pens and integrating it with patient medical records or other data collection tools remains a significant challenge.

Finally, there are a variety of papers that discuss what it takes to deploy and sustain ICTD interventions. For example, Champanis et al. [6] present deployments of a mobile application for collecting water quality data. Patel et al. [17] describe the design of an interactive voice forum for farmers in rural India, while Anokwa et al. [4] deployed a phone-based clinical decision support system in a network of hospitals in Kenya. There has also been work that focuses on strengthening health systems and CHW supply chains [19]. Our paper contributes to this understanding of what it means to put the systems in place by presenting the

deployment of a system that integrates both paper and digital technologies, and examining the impact of this system at multiple levels of the information hierarchy.

3. BACKGROUND AND CURRENT CHALLENGES

To better understand the current state of the CHW supply chain in Mozambique we conducted interviews and site visits to clarify and validate descriptions of the current program, understand the data collection procedures and identify strengths and weaknesses of the existing program.

3.1 Background

In 1978, Mozambique’s Ministry of Health began the CHW program to extend coverage of the national public health system to include underserved rural populations. During Mozambique’s civil war this program eventually became defunct and is currently undergoing revitalization. In the new program, each CHW is responsible for a population of between 500-2000 people and performs routine tasks such as developing strong ties with the community, health promotion and education, family planning counseling, and prevention and treatment of common ailments [23]. CHWs receive a monthly stipend in exchange for the services that they provide, although the stipend is not a full-time salary and the CHWs often have to balance their health worker duties with other work. To support the curative and preventative services they provide, CHWs receive a monthly kit of health supplies, including essential medicines, rapid diagnostic tests for malaria and male condoms. Each CHW is given medicines and supplies designed to treat 250 patients per month, though depending on their catchment area and the season, they may treat more or fewer. Supplies reach the CHWs through Mozambique’s network of public pharmacies. District pharmacies obtain the supply kits from the provincial pharmacy and distribute them to health centers on a monthly basis. CHWs then visit their assigned health center once a month to receive their kit of supplies.

Each district operating the CHW program has one Ministry of Health staff member, known as the district supervisor, dedicated to supervising approximately 25 CHWs in the district. The tasks assigned to the district supervisor include training new CHWs, conducting regular supervision of the CHW program, assessing the availability of medicines and supplies, and monitoring the quality of case management and proper completion of reports.

3.2 Current Challenges

Our exploration of the current CHW supply chain revealed several key challenges faced by the CHW program. One of the biggest challenges is that there is currently no standardized resupply process or system for reporting of logistics data [3]. For example, the quantity of commodities included in the CHW supply kit is based on forecasts made at a central level according to predicted disease burden and estimated monthly consumption of each commodity. However, when asked about the monthly consumption patterns of the commodities in their supply kits, different CHWs reported varying rates of consumption. This situation highlights that differences in consumption patterns between CHWs often result in CHWs either stocking out or being overstocked. However, there is currently no reporting mechanism that allows supervisors to monitor stock status of CHWs or national level partners to validate kit quantities.

In addition, although supplies are ordered regularly, they are not necessarily delivered regularly, which may result in CHWs experiencing stock-outs. However, there is no process for

documenting the frequency or severity of these stock-outs. This results in CHWs not having supplies and district supervisors not knowing the stock status of the CHWs. To address this issue, several district supervisors and pharmacy staff expressed a desire for higher quality commodity consumption data from CHWs. In addition, CHWs expressed that it is challenging for them to track their own health supplies each month. Many CHWs in Mozambique have only attended primary school and so find it difficult to understand logistics concepts and track their usage of supplies over time.

These challenges suggest that it would be beneficial to design a standardized data collection and reporting system to allow CHWs to easily collect and share logistics data with their supervisors, and make it easy for supervisors to aggregate and analyze the data and make it accessible at higher levels of the information hierarchy. Moreover, it is important to do this quickly to maximize the time available to respond to stock-outs or over-stocks.

4. INTERVENTION DESIGN

In response to the challenges identified above, we designed an intervention to introduce a standardized system for collecting and reporting CHW commodity consumption data. Our design focuses on developing a workflow that will be cheap, familiar and easy-to-use for CHWs and that will also provide district supervisors with a way to digitize, analyze and disseminate CHW logistics data.

Figure 2 shows the intervention workflow. To track CHW usage of health commodities, practitioners at the nonprofit organization VillageReach [22], as a subcontractor to the USAID | DELIVER PROJECT Task Orders 4 and 7 [2], designed a paper form that will be filled out by CHWs. The choice to use a paper form at the CHW level was made for several reasons. First, paper forms are familiar to the CHWs and well-suited to their education level and prior experience. In addition, paper forms are cheap to produce and each form can store and display a relatively large amount of data, such as the number of each of 16 commodities consumed. Furthermore, the paper form will serve as a visual and tangible record that will allow CHWs to see their supply usage from day to day. Finally, in this scenario, the paper forms will be easy to distribute because the CHWs already visit their health center supervisor on a monthly basis to collect supplies.

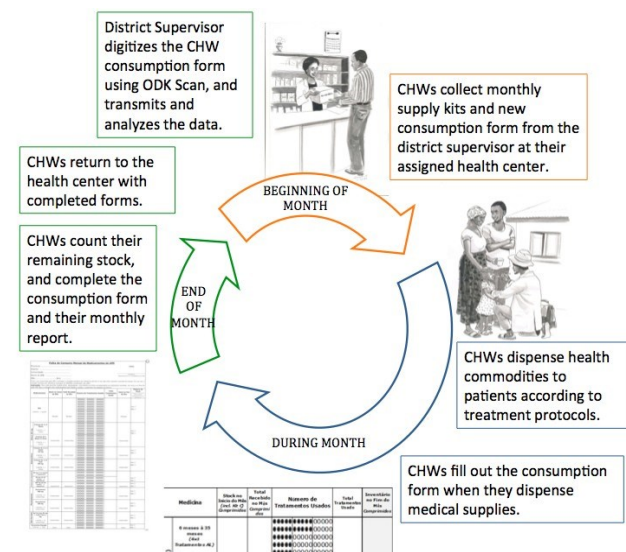


Figure 2: The monthly workflow for recording and digitizing CHW commodity consumption data [3].




Figure 3: The ODK Scan application workflow.

To complete the form, CHWs record their initial stock levels at the beginning of the month and any extra stock received during the month. They then fill in one bubble on the form for each treatment provided to patients. In addition, CHWs will use the form to record their remaining stock at the end of the month and whether they faced stock-outs during the month. They will then present the completed form to their health center supervisor when they go to pick up the next month’s supply kit. The supervisor will use ODK Scan [9] (described in detail in the next section) to digitize, aggregate and analyze the CHW consumption data and transmit the data to a cloud-based server running ODK Aggregate [12]. The data can then be used by the Ministry of Health or partner organizations to keep track of CHW activities over time and to recognize when CHWs experience stock problems so that they can be provided with supplementary stock if necessary.

To summarize, the intended outcome of this intervention is that CHWs will use the paper form to track their stock on hand, quantity received, quantity dispensed (consumption), and stock-outs, and regularly report this data to their district health supervisor, who will digitize and transmit the data using ODK Scan so that it is immediately available for all stakeholders.

5. OVERVIEW OF ODK SCAN

ODK Scan is an Android application that uses computer vision techniques to digitize data from paper forms. Many government, social and health organizations in developing countries rely heavily on paper forms to perform large-scale data collection. Paper forms are a well-understood and trusted medium, and the low cost and ease-of-use of paper suggest that it will continue to be extensively utilized for many years to come. However, the benefits of digitizing data from paper forms for the purposes of statistical analysis and aggregation are significant, and providing stakeholders with quick access to accurate data could aid critical decision making regarding resource allocation, assessment and planning. Many paper forms used by organizations for data collection contain a mixture of data types, including handwritten text, numbers, checkboxes and tallies in “bubble” format. While some of these data types, like handwritten text, require a person to transcribe the data, others, like checkboxes or bubbles, can be interpreted by a computer.

ODK Scan is capable of automatically classifying “machine-readable” data, like bubbles and checkboxes, and saves image “snippets” of form fields that are not machine-readable. These snippets are incorporated into the ODK Collect [12] interface to aid the entry of non-machine readable data like text and numbers. ODK Scan uses a JSON [8] form description language to facilitate the processing of existing paper forms without the need to

redesign the forms. The ODK Scan workflow is shown in Figure 3. To digitize a form, a user captures an image of the form using the standard Android camera app. The captured image is then passed to ODK Scan where it is aligned using a template image of an empty form. After alignment, the form fields are extracted and machine-readable data types are processed automatically. Users are then prompted to enter values for non-machine readable fields, which are displayed one at a time in ODK Collect with an image snippet of the field. In addition, since ODK Scan occasionally makes classification errors, the application allows users to check and possibly correct the results of the automated scanning.

There are two prior research papers that describe different aspects of the ODK Scan system. The first [9] describes the technical design and initial smartphone implementation of the ODK Scan algorithms and shows that the software is capable of processing “bubble” fields with over 99% accuracy. The second [10] presents the design and laboratory evaluation of user interfaces and input techniques to aid the entry of non-machine readable data using image snippets. Together, these papers establish that ODK Scan is capable of digitizing data from paper forms with a high level of accuracy, and also that data entry using ODK Scan and snippet-enabled ODK Collect is more efficient than manual data entry. However, the papers did not evaluate a real deployment of the application, and did not consider the benefits, challenges and limitations associated with deploying and sustaining the application in a low-resource setting.

Therefore, this paper focuses on the integration of ODK Scan into a real information ecosystem: the health worker supply chain in Mozambique. From a development perspective, there were several changes that had to be made to facilitate the scanning of CHW supply consumption forms. Many of these changes are likely to also be beneficial for other users of the application. For example, the large number of CHW supplies and other data that needed to be collected from health workers resulted in the form being two pages instead of a single page. As a result, the application was modified to correctly handle forms that have a variable number of pages. In addition, since several data values on the consumption form often remain the same from month to month, ODK Scan was modified to incorporate “default” values for fields. Form fields will be pre-populated with the specified default value and would only need to be entered if the value on the paper form differs from the default value.

6. IMPLEMENTATION

This section describes the work required to incorporate ODK Scan as part of a new system for collecting and reporting of CHW commodity consumption data in Mozambique. We focus on

documenting the experiences of three primary stakeholders: the NGO practitioners responsible for implementing the technology, the CHWs who collect and record data on paper forms in the field, and the district supervisors who use ODK Scan to digitize the collected data.

6.1 Practitioners

6.1.1 Form Design and Description

The first step in the implementation was to design an ODK Scan compatible paper form to be used by CHWs to record the commodities that they consume. The form, shown in Figure 4, was designed by practitioner staff at the USAID | DELIVER PROJECT with input from district and provincial health staff and other partner organizations. The form design incorporated several features to facilitate digitization by ODK Scan. For example, the form contained three pictures located in the corners of the form to assist ODK Scan with feature extraction for form alignment. In addition, the form fields that record the number of supplies consumed consist of machine-readable bubble tallies that can be automatically processed by the application.

In addition to designing the paper form, practitioners at the USAID | DELIVER PROJECT also created the corresponding ODK Scan form description file to tell the application the name, size and data type of each form field. Creating this form description file was the most technically challenging aspect of the entire deployment. To make it easier to specify the form fields, we developed a JavaScript web application called the “template maker” that allows users to drag and drop a form image into the browser window and mark up different form fields by drawing boxes on the image. The software then automatically creates a JSON object that describes the size and location of the boxes. To specify a bubble, the user clicks on a point within an already defined form field. If the location of the bubble is not exactly correct, the user may refine the bubble coordinates by editing the automatically generated JSON in a side panel on the user interface. All newly created fields are assigned a default “bubble” classifier. By editing the JSON that specifies the classifier, users may customize the size or type of the bubbles on the form.

Creating the form description file required significant communication between the ODK Scan developers and the NGO practitioners designing the form. Over several rounds of iterative design, the developers worked to make the template maker easier to use and to add features requested by the practitioners, and the practitioners used the tool to create a description file for the form. The challenges faced and lessons learned from this process and the design decisions that resulted from the inherent technical nature of the form descriptions are discussed in Section 9.2.

6.1.2 Training Materials and Documentation

In low-resource settings, many practitioners or end-users have limited experience with technology. As a result, it is crucial to develop comprehensive documentation and training materials that do not assume prior knowledge of the application or device. To aid other practitioners who want to setup, configure and deploy ODK Scan we developed a “setup and technical guide”. This guide provides instructions on basic functionality such as how to power up and navigate the device and access the SD card. It then covers how to install and configure ODK Scan, Collect and Aggregate. Finally, the guide covers how to handle a variety of common mistakes and error messages. We also created a user manual for district supervisors that demonstrates how to place the phone’s camera above the form and take good quality pictures. Step-by-step instructions and screenshots guide the user through

Figure 4: The two-page ODK Scan compatible CHW health commodity consumption form.

capturing and processing form images using ODK Scan, entering non-machine readable data using image snippets in ODK Collect, and transmitting data to ODK Aggregate. The manual also addresses common errors or problems. Finally, we also designed a CHW help booklet with pictures and instructions on how to correctly fill out the paper consumption forms.

6.2 CHWs

In addition to providing CHWs with help booklets, we also held a one day training workshop (from 8am-7pm) in each district to teach CHWs how to correctly use the consumption forms. The district supervisor and approximately 25 CHWs from each district participated in each workshop. The day began with classroom sessions that explained how to fill out the forms, after which CHWs completed a series of practical exercises in small groups that included a large amount of practice filling out forms. Since many CHWs in Mozambique only have primary school education, the groups were specifically designed to allow better educated CHWs to help less educated CHWs understand the concepts. To evaluate the efficacy of the workshop, CHWs completed a short quiz before and after the training sessions that assessed their knowledge of how to use the forms. CHWs who had trouble with the quiz received additional help until they demonstrated sufficient confidence with the form. In addition, since tracking commodities and recording stock-outs were new concepts for many CHWs, we also conducted a half-day follow-up training session one month after the initial training that provided CHWs with additional practice and assessment on their filled out forms.

6.3 District Supervisors

In addition to the CHW training workshops, each district supervisor received a full-day one-on-one training session on how to use the ODK tools to digitize forms and transmit data. To begin, a field staff member from the USAID | DELIVER PROJECT demonstrated how to digitize several consumption forms. Then the district supervisor digitized approximately 15 forms under the supervision of the field staff member. In addition to using the software, the supervisor also received advice on how to maintain the phone and keep it secure, and on how to assess the

quality of form completion and give advice to CHWs if the forms were not filled out satisfactorily.

7. QUANTITATIVE EVALUATION

To evaluate the impact of integrating ODK Scan into the CHW supply chain, we conducted a four month study in two districts in Maputo province: the Manhiça and Marracuene districts. A total of 45 CHWs (22 from Manhiça and 23 from Marracuene) were trained to fill out the consumption forms and two district supervisors were trained to digitize the forms using ODK Scan. Training was completed in December 2012 using the methods and training manuals described in the previous section. After the training sessions, the application was deployed for a period of four months between January and April 2013.

7.1 Data Collection

To measure the amount of CHW consumption data collected, we counted the number of forms that CHWs submitted to district supervisors and the number of forms that the supervisors digitized using ODK Scan, Collect and Aggregate. We also evaluated the quality of forms with respect to folding, warping and cleanliness.

7.1.1 Results

Over the four month study period, the 45 CHWs submitted a total of 140 paper forms (out of a possible 180) to the district supervisors. There were a variety of reasons why CHWs sometimes did not submit forms. Several lived very far away and were unable to submit forms because of transport difficulties. Others did not fill out forms because they did not receive their monthly kit of supplies. Some were simply not working or had left the area. Of the 140 forms received, district supervisors scanned, processed and uploaded a total of 122 forms. The supervisors said that the remaining 18 forms were not scanned because they contained too many errors to yield useful data. A large proportion of the forms not scanned were from February 2013, and subsequent investigation revealed that many CHWs did not fill out the form correctly at the beginning of the month because they had not received their monthly stipend from the Ministry of Health. In addition, many of the 122 forms that were digitized and transmitted were in poor condition: dirty, warped or folded. These artifacts contributed to a higher than expected ODK Scan error rate that we will now discuss in detail.

7.2 Accuracy

Each form contains 6 text fields, 68 numeric fields, 21 yes/no checkbox fields and 1395 bubbles in 16 tally fields. To measure the accuracy of the data collected, we selected a random sample of 35 consumption forms and manually created groundtruth data for these forms. Then we evaluated the accuracy of ODK Scan by comparing ODK Scan's bubble and checkbox classifications to the groundtruth data. We evaluated the combined accuracy of ODK Scan and Collect by comparing the data submitted to ODK Aggregate with the groundtruth data.

To tease apart the proportion of scanning errors that occurred as a result of the ODK Scan algorithms and the proportion that occurred due to errors in the form description or the poor condition of the forms, we asked an expert ODK Scan developer to create a new form description file for the consumption form and compared the results obtained using the expert's form description with the results obtained using the original NGO practitioner's form description. The expert's form description also enabled several additional form alignment features. These features were not enabled in the deployment because initial testing of the form design revealed that they had a slightly negative effect on the

Table 1: Bubble classification accuracy on a sample of 35 consumption forms using the practitioner's form description and the expert's form description.

	Practitioner's Form	Expert's Form
True Negatives	39387	42699
False Negatives	867	354
True Positives	6537	7050
False Positives	3504	192
Overall Accuracy	91.3%	98.9%

overall classification accuracy. The forms that we used to test the form design were in relatively good condition compared to the real consumption forms, and the extra alignment features appeared to slightly overcorrect the form alignment when corrections were not necessary. However, since many of the real consumption forms were in very poor condition, we wanted to see how the extra alignment features might affect the accuracy of processing.

7.2.1 Results

Table 1 shows the results of comparing the automated ODK Scan classifications to the groundtruth data. Overall, we saw a classification accuracy of 91.3% with the practitioner's form description and 98.9% with the expert's form description and extra alignment features. A large number of the errors that occurred with the practitioner's form description resulted from false positive classifications and a closer examination of the form description revealed that many specified bubble locations were in fact a few pixels away from the true center of the bubble. Thus, about half of the improvement in accuracy that we saw with the expert's form description resulted from correctly specified bubble locations and the other half resulted from extra alignment features.

To assess the impact that the incorrectly classified bubbles had on the number of fields that required correction in ODK Collect, we calculated the percentage of fields in which the ODK Scan tally differed from the groundtruth tally. Using the practitioner's form description, 39.4% of fields differed from the groundtruth, whereas with the expert's form description and extra alignment features, 12.3% of fields differed from the groundtruth. In both cases, the values of most of these fields differed from the groundtruth values by a small number of bubbles, as shown in Figure 5. The few cases where the ODK Scan values differed from the groundtruth by a large number (more than 10 bubbles) was due to alignment errors in which an entire form field was misaligned. Fortunately, this kind of error occurred infrequently.

Despite the relatively high error rate that occurred with the practitioner's form description, we observed a high level of overall accuracy in the data that the district supervisors submitted to ODK Aggregate. The accuracy of manually entered numeric and text fields was 98.5%, and the overall accuracy of all data types submitted was 98.1%. This shows that users were correcting the errors made by ODK Scan. The implications of the differences in accuracy between the practitioner's form description and the expert's form description are discussed in Section 8.3.1.

7.3 Time

To measure how long it took to digitize data, we timed the district supervisors as they digitized several paper forms every month. We also separately measured how long it took for the supervisors to complete the ODK Scan portion of the process followed by ODK




Figure 5: Percentage of form fields with differences between ODK Scan tallies and groundtruth tallies using the practitioner's form description and the expert's form description with additional alignment features.

Collect with image snippets. On average it took district supervisors 12 minutes to complete the scanning and data entry for each two-page consumption form. The ODK Scan portion took about 2 minutes and the data collection in ODK Collect took an additional 10 minutes. Interestingly, the amount of time required for data entry did not change substantially from month to month, with supervisors needing about the same amount of time to digitize data at the end of the study as at the beginning. Feedback obtained from the supervisors and reasons for the amount of time required to digitize the data are discussed further in Section 9.2.

7.4 Cost

We assessed the total cost of deploying and sustaining the intervention by summing the total costs of the devices, paper forms, data transmission, data storage and personnel. The two district supervisors in the study were each issued an HTC Nexus One Android device that cost \$320. Obtaining and printing all the paper forms cost less than \$25 per month. The airtime for transmitting the digitized data cost \$0.05 per form (or roughly \$1.50 per month) and storing the data in the cloud cost \$8.50 per month. Additional costs included the transportation and per diem costs for the 45 CHWs and 2 district supervisors to attend the training sessions, and the salaries of the practitioner staff who implemented and supported the deployment. We estimate that monitoring the CHWs and district supervisors required 10-15 days worth of field staff time per month, while designing the form and creating the form description required 10 days of technical staff time. Thus, as is common in ICTD interventions, we found that personnel costs far exceeded the costs associated with the technology and we recommend that organizations pay close attention to this factor for any future deployment.

8. QUALITATIVE EVALUATION

To understand the benefits, challenges and limitations associated with deploying and sustaining the intervention we collected and analyzed qualitative data from three different user groups: NGO practitioners, CHWs and district supervisors. Prior studies have shown that qualitative responses may suffer from participant response bias [11, 13]. To mitigate the potential for response bias, we conducted multiple interviews and observations over the entire study period and compared how users' experiences and attitudes changed over time. In addition, to elicit varied responses from participants, the interviews at the end of the study used different

language and wording to the interviews conducted during the study. Finally, we also compared the qualitative data obtained between user groups to better understand the impact across the different levels of the information hierarchy.

8.1 Practitioner Experiences

We conducted several semi-structured interviews with each of the practitioner staff involved in the intervention. Each interview lasted approximately one hour and involved discussing in detail all of the tasks performed by the practitioner, their opinions and experiences with the technology, and any additional challenges or issues that arose while supporting the CHWs and district supervisors in the field.

8.1.1 Findings

Practitioners reported that designing the consumption form was time-consuming because many stakeholders (NGO and Ministry of Health staff) wanted different data points included on the form. This resulted in a large amount of design iteration, and the form layout and fields changed every time the design changed. In addition, each new design had to go through a review and approvals process and a new form description had to be created.

Creating the ODK Scan form description was also time consuming and challenging. Although the template-maker software provided a GUI to reduce the amount of JSON that had to be written, the form designer still had to manually edit JSON to adjust bubble locations and provide labels and identifiers for form fields. The form contained 21 checkboxes and 1395 bubbles in 16 tally fields. To specify a bubble location, the form designer clicked on the center of the bubble in the template maker. However, it was easy to miss the center by a few pixels, and adjusting the bubble location required manually editing coordinates in a side panel on the user interface. In addition, every time the form changed, a large amount of work was required to adjust the form description, including updating the locations of each bubble. It is likely that the large number of changes to the form design and description introduced some systematic alignment issues that led to the relatively high error rate encountered in the field. This is discussed further in Section 9.2.

Practitioner field staff also reported facing several challenges while supporting CHWs in the field. Many CHWs lived and worked in very remote places, and in some cases it took over 4 hours to reach a CHW. On several occasions the practitioner was unable to reach a CHW because the road was washed away. An important implication of this finding is that it must have also been challenging or impossible for the CHW to travel to the health center to receive a supply kit.

8.2 CHW Experiences

To evaluate the experiences of the CHWs, we chose a random sample of 8-9 CHWs each month (4-5 from each district) and visited each CHW in their community to observe and assess their performance. In total, we visited 33 CHWs over the four month study period. Each visit lasted between 30 minutes and 2 hours. To evaluate the accuracy of the CHW's consumption form, we checked their stock level for each health commodity and compared their real stock level with the stock level on the form. We also checked the CHW's patient register, calculated the number of each commodity that had been distributed and compared that number to the amount indicated on the form. Finally, we performed a semi-structured interview with the CHW to understand their experiences with the consumption form and other issues that may have affected their performance.

8.2.1 Findings

The majority of the CHWs visited were correctly using the form to record their consumption of medicines, and the amount of stock they possessed matched the stock recorded on the form. However, several CHWs experienced problems with the form. The most common problem was that CHWs would forget to record their initial stock levels at the beginning of the month and would later try to guess how much stock they had started with. This problem occurred with 7 of the 33 CHWs that we visited and was exacerbated in one district during the month of February because the CHWs did not receive their monthly stipends from the Ministry of Health on time and so lacked the motivation to record their consumption data. For 2 of the 33 CHW visits we found the data on the consumption form did not match the physical stock levels or patient register, and for another 2 we found the CHWs were not using the form because it was too complicated and they did not know how to fill it correctly.

There were also several occasions where delivery delays resulted in CHWs not receiving their full kit of health supplies on time. As a result, some CHWs were forced to improvise to try and meet the needs of their patients. For example, when some CHWs did not receive their stock of 250mg paracetamol tablets, they would sometimes break 500mg tablets in half and distribute them as 250mg tablets. This behavior affected the data on the forms since the CHWs reported distributing 250mg tablets despite not having received any stock that month. In addition, one of the medicines was delivered to CHWs in packages of nine tablets. As a result, to record the number of tablets distributed, CHWs had to multiply the number of packages by nine. Since many CHWs have only attended primary school, several of them found it difficult to do this calculation correctly and made mistakes on the consumption forms. We also observed some interesting additions made to forms. For example, as shown in Figure 6, one CHW ran out of bubbles to fill in and improvised by drawing additional bubbles on the form to record the commodities distributed to patients.

8.3 District Supervisor Experiences

To assess the experiences of the district supervisors, we visited each supervisor once a month for the four months of the study. Each visit lasted several hours, in which we checked on the state of the phone and watched the supervisors digitize data from several consumption forms. We also conducted a semi-structured interview to monitor the supervisors' experiences with the application and document any issues that came up during the month. At the end of the study, we performed an additional semi-structured interview with each district supervisor to further understand the impact that the new technology had on their work.

8.3.1 Findings

At the start of the study, each district supervisor reported owning a feature phone, although neither had used a touchscreen device before. As a result, they were somewhat hesitant to use the phone at first and expressed concern that they might break the phone or the application. However, their proficiency with the application increased substantially after the first month and they became visibly more confident and comfortable with the technology. Both




Figure 6: Excerpt of a consumption form showing additional bubbles drawn in by the CHW.

supervisors were careful with the phones and kept them locked up when not in use. None of the phones were broken, lost or stolen during the study. The supervisors were able to charge the phones and did not experience any hardware or software malfunctions. In general, the supervisors reported sufficient network connectivity to transmit data to the server. This was supported by the fact that the server successfully received regular data transmissions from the supervisors throughout the study. However, the supervisors also said that sometimes they needed to go outside to find a spot with good connectivity before being able to transmit data. This confirms the importance of being able to store data on the phone and transmit it when a connection is available.

Supervisors described data entry with ODK Scan as being more efficient than paper based summaries because they were able to process the form from anywhere since an image of each data field was displayed on the phone. However, the supervisors also found verifying the data in ODK Collect and manually entering non-machine readable data types to be time-consuming. One reported,

“In the beginning I took the whole day or hours to do the job. I had all my focus on it, and when I stopped, I lost my concentration.”

One reason that it took a long time to enter data may be the decreased level of ODK Scan accuracy discussed in Section 7.2. Another reason could be that the form contained 68 numerical fields that needed to be typed into the phone. Form designers should therefore try to minimize the number of fields that must be manually entered and should also provide likely default values for as many fields as possible.

Although it took a long time to digitize the forms, the supervisors both felt that using the application was more efficient and saved them time over a purely paper-based reporting system. One supervisor told us,

“Using ODK Scan it is quicker to send the data to anywhere. When I have written data, I need to write everything down, then take it for [an approval] signature, then find transport to send that information to [the provincial office]. It can take days to do that.”

In addition, both supervisors noted that the forms made it easier to track health supplies in a standardized way. One supervisor said,

“Before [this intervention], we used to control the medicines in a rudimentary way because there were different pieces of information in different sources. This form provides the information all in one place. It is very useful for me. I can just see the form and know what happened. Before, I would have to go to different data sources to confirm everything.”

In addition, they appreciated no longer needing to prepare handwritten reports:

“The system facilitates the sending of reports, so there is no need for me to move.”

Finally, the supervisors cared about the accuracy of the data. We observed that they sometimes used white-out to correct errors on CHW consumption forms, and they also took time to use ODK Collect to correct classification errors made by ODK Scan. The supervisors also recognized that the data collected would be beneficial for other stakeholders:

“This information is very useful for the pharmacy to see the consumption patterns and to know how much [stock] should be sent to the district”

and they recommended that the intervention be scaled up:

“This is very important [for other districts] because we can now track the consumption of medicines. It was something that we weren’t doing before this was started. It is important for us to know and it helps us to know how to use the medicines in the right way.”

9. DISCUSSION

9.1 Study Size

The success of any intervention depends on the input and buy-in of all partners, including public health staff at all levels, NGO and technical assistance staff, and end users. Although our study targets only two district supervisors as the end-users of ODK Scan, a much larger number of people were impacted by the choice to integrate the technology into the CHW supply chain. To support the district supervisors, several practitioner staff learned to use ODK Scan, created training and trouble-shooting materials, organized training sessions, and supported the district supervisors throughout the study period. Several other technical and managerial NGO staff designed the consumption form, consulted with other stakeholders and ensured that the form was compatible with ODK Scan. Finally, 45 CHWs had to learn how to use the consumption forms. This study therefore impacted at least 2 district supervisors, 5 NGO practitioner staff and 45 CHWs. Understanding the impact that integrating ODK Scan has at this relatively small scale, but at multiple levels of the information hierarchy, is a critical step in understanding how the technology might be successfully scaled up.

9.2 Form Design and Description

It is encouraging that the NGO practitioners were able to design an ODK Scan compatible form, create the corresponding form description file and successfully deploy the application. However, creating the form description file was technically challenging and required substantial time and effort. In addition, the impact that using a slightly inferior form description had on the accuracy of scanning and subsequently on the number of fields that had to be corrected in ODK Collect suggests that generating a highly accurate form description file should be a priority for any organization wishing to use ODK Scan. In addition, rigorously testing the form description with forms that are in very poor condition is essential to ensure that the extra ODK Scan alignment features are enabled if necessary.

The biggest challenge in creating the form description file was specifying the locations of hundreds of bubbles with pixel-perfect accuracy. When specifying bubble locations by clicking on points in a form field, it was easy to miss the true center of a bubble by a few pixels, which led to a higher number of incorrect bubble classifications. In addition, whenever the form layout changed slightly, a large amount of work was required to also change the form description. Although we added several transformation controls to address this problem, we found that generally the bubble locations did not match up perfectly after a transformation due to rounding errors and other minor form differences. Finally, the current version of the software requires that the user know how to write and edit JSON. A more polished GUI that hides the JSON from the user might have decreased the amount of JSON that had to be written, although we saw that the ability to copy and paste sections of JSON also made editing form descriptions easier.

To make it easier for users to create accurate form descriptions, we are building a new software tool that will automatically generate an ODK Scan compatible form image and form description file from an Excel spreadsheet. Users will specify

form questions in Excel and all of the required ODK Scan files will be generated automatically. Although this will not make it easier to add existing forms to the system, we believe that it will lower the barrier to use and allow more people to quickly and accurately define ODK Scan compatible forms.

9.3 Lessons Learned and Recommendations

Many published ICTD deployments are conducted by developers or researchers, not by practitioners. Although practitioners are frequently able to imagine the potential benefits afforded by a new technology, they often do not understand the technical expertise or effort required to deploy the technology successfully, which may lead to false expectations. One of the major lessons learned in this work is that it required a substantial amount of time, effort and patience on the part of the practitioners to successfully deploy and sustain the intervention. Practitioners had to learn how to use the technology themselves and also how to provide support to CHWs and supervisors through regular visits and monitoring. Without this effort it is unlikely that the deployment would have been successful.

One of the major benefits of ODK Scan is that it allows people to continue to use cheap and familiar paper forms at the bottom level of the information hierarchy. However, the ease of use of paper forms also means that people are likely to use them in unexpected ways. For example, in this study, we observed district supervisors using white-out to correct errors on consumption forms prior to using ODK Scan. We also saw CHWs sometimes draw additional bubbles on the forms (see Figure 6). These findings highlight an interesting trade-off between human-readable and machine-readable data. On one hand, it is clearly beneficial to have a person verify all of the processed data values. On the other hand, requiring that a person verify all of the data increases the time required for data entry. Ultimately it will be up to the user to manage this trade-off and decide the level of human verification required for a given application. An important recommendation for organizations who want to use ODK Scan is to make sure that the paper form contains as many machine-readable fields as possible. In this study, the CHW consumption form contained a large number of numeric fields. Since ODK Scan cannot automatically process handwritten numbers or text, these fields had to be manually transcribed which was time-consuming.

We also experienced a number of implementation issues that affected the deployment. For example, district supervisors did not always receive consumption forms from some CHWs for a variety of reasons: some CHWs lived very far away and were unable to submit forms, some did not fill out the forms because they did not receive their kit of supplies, while others were simply not working or had left the area for personal reasons. When asked about the missing forms, one supervisor told us, *“The program will only work as much as the people work”*. In addition, as noted by another supervisor, *“In February we had a lot of mistakes [on the forms] due to lack of CHW motivation because they did not receive their stipends.”* However, despite these issues the supervisors were happy to receive data from the CHWs that did submit forms and frequently stated that having access to the data improved their knowledge and decision-making ability regarding the CHWs’ medical commodity needs, and also helped them to calculate the resupply amounts required.

10. FUTURE WORK AND CONCLUSION

We have presented an in-depth analysis of the work required to integrate ODK Scan into the CHW supply chain in Mozambique. ODK Scan was successfully deployed in two districts for four months and is still deployed at this time. The application has proved

to be a usable and useful tool for recording CHW commodity consumption data. We have already obtained permission to deploy the intervention in four more health districts in Mozambique and we expect this number to expand. The lack of consumption data prior to this intervention will make it interesting to see if having access to the data changes existing CHW monitoring and resupply processes.

We have found the process of identifying and documenting the benefits and challenges associated with deploying and sustaining ODK Scan to be extremely valuable and have learned many lessons that we hope will be useful to other researchers and practitioners who want to deploy ODK Scan or other applications that bridge the gap between paper-based and digital data collection in low-resource settings. Finally, we have specifically highlighted the crucial role played by practitioners in this deployment. We hope that our approach will inspire more work that focuses on strengthening the link between researchers and practitioners in ICTD.

11. ACKNOWLEDGEMENTS

We would like to acknowledge the Mozambique Ministry of Health and all our study partners and participants. This work was funded by a grant from the Bill & Melinda Gates Foundation through the Grand Challenges Explorations Initiative, the USAID | DELIVER PROJECT and National Science Foundation grant 1111433.

12. REFERENCES

- [1] CORE Group, Save the Children, BASICS, and MCHIP. 2010. *Community Case Management Essentials: Treating Common Childhood Illnesses in the Community: A Guide for Program Managers*. Washington, DC, CORE Group.
- [2] USAID | DELIVER PROJECT. <http://deliver.jsi.com>.
- [3] USAID | DELIVER PROJECT, Task Order 4. 2012. *Mozambique: Strengthening the Community Health Worker Supply Chain*. Arlington, VA.
- [4] Y. Anokwa, N. Ribeka, T. Parikh, G. Borriello, and M. Were. Design of a Phone-Based Clinical Decision Support System for Resource-Limited Settings. In *International Conference on Information and Communication Technologies and Development (ICTD)*, 2012.
- [5] Captricity. <http://captricity.com/>.
- [6] M. Champanis and U. Rivett. Reporting Water Quality A Case-study of a Mobile Phone Application for Collecting Data in Developing Countries. In *International Conference on Information and Communication Technologies and Development (ICTD)*, 2012.
- [7] K. Chen, A. Kannan, Y. Yano, J. Hellerstein, and T. Parikh. Shreddr: pipelined paper digitization for low-resource organizations. In *ACM Conference on Computing for Development (DEV)*, 2012.
- [8] D. Crockford. Introducing JSON. json.org, 2009.
- [9] N. Dell, N. Breit, T. Chaluco, J. Crawford, and G. Borriello. Digitizing paper forms with mobile imaging technologies. In *ACM Conference on Computing for Development*, 2012.
- [10] N. Dell, N. Breit, J. Wobbrock, and G. Borriello. Improving Form-Based Data Entry with Image Snippets. In *ACM Graphics Interface Conference (GI)*, 2013.
- [11] N. Dell, V. Vaidyanathan, I. Medhi, E. Cutrell, and W. Thies. "Yours is better!" Participant Response Bias in HCI. In *ACM Conference on Human Factors in Computing Systems (CHI)*, 2012.
- [12] C. Hartung, Y. Anokwa, W. Brunette, A. Lerer, C. Tseng, and G. Borriello. Open Data Kit: Building Information Services for Developing Regions. In *International Conference on Information and Communication Technologies and Development (ICTD)*, 2010.
- [13] J. Ledlie. Huzzah for my Thing: Evaluating a Pilot of a Mobile Service in Kenya. *Qual Meets Quant*, ICTD, 2010.
- [14] I. Medhi, A. Tewari, M. Jain, and E. Cutrell. The Fate of a Digital Slate: Unexpected Issues with Deployment in Rural India. *User Experience Magazine*, 11(2), 2012.
- [15] G. Mhila, B. DeRenzi, C. Mushi, T. Wakabi, M. Steele, P. Dhaldialla, D. Roos, C. Sims, J. Jackson, and N. Lesh. Using Mobile Applications for Community-based Social Support for Chronic Patients. *HELINA*, 2009.
- [16] T. Parikh, P. Javid, K. Sasikumar, K. Dhosh, and K. Toyama. Mobile Phones and Paper Documents: Evaluating a New Approach for Capturing Microfinance Data in Rural India. In *ACM Conference on Human factors in computing systems (CHI)*, 2006.
- [17] N. Patel, D. Chittamuru, A. Jain, P. Dave, and T. Parikh. Avaaj Otalo A Field Study of an Interactive Voice Forum for Small Farmers in Rural India. In *ACM Conference on Human Factors in Computing Systems (CHI)*, 2010.
- [18] A. Ratan, S. Chakraborty, P. Chitnis, K. Toyama, K. Ooi, M. Phiong, and M. Koenig. Managing Microfinance with Paper, Pen and Digital Slate. In *International Conference on Information and Communication Technologies and Development (ICTD)*, 2010.
- [19] H. Schneider, D. Blaauw, L. Gilson, N. Chabikuli, and J. Goudge. Health Systems and Access to Antiretroviral Drugs for HIV in Southern Africa: Service Delivery and Human Resources Challenges. *Reproductive Health Matters*, 14(27):12 – 23, 2006.
- [20] H. Underwood, R. Sterling, and J. Bennett. The Design and Implementation of the PartoPen Maternal Health Monitoring System. In *ACM Conference on Computing for Development (DEV)*, 2013.
- [21] S. Van Wart, K. Tsai, and T. Parikh. Local Ground: A Paper-Based Toolkit for Documenting Local Geo-spatial Knowledge. In *ACM Conference on Computing for Development (DEV)*, 2010.
- [22] VillageReach. <http://villagereach.org/>.
- [23] WHO. Global Experience of Community Health Workers for Delivery of Health Related Millennium Development Goals, 2011.