Visible Communities: Designing a Socio-Spatial Map

Raphael Schaad
B.S., Bern University of Applied Sciences (2011)

Submitted to the Program in Media Arts and Sciences, School of Architecture and Planning, in partial fulfillment of the requirements for the degree of

Master of Science in Media Arts and Sciences

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY (MIT)

June 2017

Author ....
Program in Media Arts and Sciences, MIT Media Lab
May 12, 2017

Certified by ....
Deb Roy
Associate Professor of Media Arts and Sciences, MIT Media Lab
Thesis Supervisor

Accepted by ....
Pattie Maes
Professor of Media Technology, MIT Media Lab
Academic Head, Program in Media Arts and Sciences

Visible Communities:
Designing a Socio-Spatial Map

Raphael Schaad

Submitted to the Program in Media Arts and Sciences, MIT School of Architecture and Planning, on May 12, 2017 in partial fulfillment of the requirements for the degree of Master of Science in Media Arts and Sciences.

Abstract

This thesis presents a collaborative human–machine crowdmapping approach to creating *socio-spatial maps* that represent both spatial and social aspects of communities. Our implemented system combines satellite image analytics, a mobile mapping app, and social survey data. The system is designed to provide an end user experience that aligns institutional interests with grassroots interests, resulting in a self-sustaining system. In collaboration with the global health organization Partners in Health, we tested our approach with local health workers in Rwanda. Better maps can improve local visibility and empower communities to share knowledge, trade goods, and access medical services. Assisted by automatically annotated satellite maps, the community-driven mapping resulted in detailed spatial and social maps for four rural villages. With the collected data, we designed a novel socio-spatial map for this community that combines knowledge about household locations, paths, inhabitants of those homes, and social relations between residents. Generalizing from this map, we propose a framework to organize *people, places, paths,* and *relationships* to reason about the intersection of social and spatial mapping. Furthermore, we derive design characteristics of our human–machine mapping system that can guide the development of new systems in related contexts. Socio-spatial maps have the potential to be used as critical decision-making tools for individuals and organizations alike.

Thesis Supervisor: Deb Roy
Title: Associate Professor of Media Arts and Sciences, MIT Media Lab

Visible Communities: Designing a Socio-Spatial Map

Raphael Schaad

The following served as readers for this thesis:

Thesis Reader . . . .

Ethan Zuckerman
Associate Professor of the Practice in Media Arts and Sciences
MIT Media Lab

Thesis Reader . . . .

Ophelia Dahl
Co-founder, Chair of the Board
Partners in Health

Acknowledgments

I express my sincere acknowledgments to:

community health workers in Burera for participating in the field study – the dedication to your communities is inspiring;

my research advisor Deb Roy for giving me the opportunity to explore socio-spatial maps, shaping my thinking, and mentoring me along this journey;

thesis reader Ophelia Dahl for teaching me about community health programs, and her colleagues at Partners in Health in Rwanda for supporting the field work: Alex Coutinho, Fredrick Kateera, Fabien Munyaneza, and dozens of other staff that supported this project;

thesis reader Ethan Zuckerman for valuable advice on this thesis, and his colleague at the MIT Center for Civic Media, Emilie Reiser, for guidance on community engagement;

senior collaborators of my research group, who all influenced this work: Andrew Heyward, James Kondo, Mike Koehrsen, Preeta Bansal, Russell Stevens, William Powers, and the most supportive lab admin, Heather Pierce;

postdoctoral researcher Soroush Vosoughi and graduate students Nabeel Gillani and Martin Saveski for introducing me to social network analysis, Anneli Hershman for advice on survey design, and all graduate students at the Laboratory for Social Machines for being great colleagues: Ann, Eric, Ivan, Juliana, Luke, Mina, Prashanth, Saquib, and Sneha;

the MIT Tata Center for Technology and Design for supporting the field research with funding and mentorship, especially postdoctoral researcher Selma Duhović;

employees at Google and Alphabet for flying the satellites and providing advice on machine learning and design sprints: Matt Hancher, Kristi Bohl, Noel Gorelick, and John Zeratsky;

Marcos Weskamp for being a mentor and inspiring designer, and John Maeda for introducing me to the magic of the Media Lab;

Eugene Yi and Catherine Tsai for editing my writing and shaping my thinking significantly;

my friends and family for their love and continued support through my years at MIT;

Your wisdom and kindness helped this work become a reality. Thank you.

Raphael Schaad, Cambridge Massachusetts, May 2017

Every good story starts with a map ...

Contents

Figures

Tables

Abbreviations

MOH	Ministry of Health
BOP	base of the socio-economic pyramid
GPS	Global Positioning System
PIH	Partners in Health
CHW	community health worker
GIS	geographical information system
UI	user interface

1 Introduction

1.1 The Digitally Invisible

Maps are important decision-making tools. We use geographical maps to situate ourselves in our environment and navigate within it. Maps also help organizations to plan and allocate resources. With the advent of digital tools, the most complete maps are stored electronically, and we access them through our computers and smartphones. However, vast regions of the world are unmapped, and communities living there are digitally invisible. Maps at the house- and road-level of the entire world are expensive to create because mapping requires specialized tools and is laborious. This data gap in mapping exists especially in low-income economies because major map services lack commercial incentives, governments are resource-constrained, and digital self-mapping tools for these low-income communities with lower technical literacy do not exist (Corbet 2009, p. 17).

Better maps can improve local visibility and empower communities to share knowledge, trade goods, and access medical services. Specifically, this thesis considers the potential of mapping and navigation services for frontline health workers in rural areas. The job of these workers requires constant travel to visit households across an assigned catchment area to provide basic health services. They might be expected to cover hundreds of homes in their catchment, ideally optimized to visit those with the acutest needs in the most expeditious manner possible. Thus, detailed knowledge of homes, paths, and medical infrastructure could significantly improve their ability to provide services. From an institutional point of view, accurate knowledge of the count and location of people and their homes would be valuable inputs for optimally allocating resources to areas of need.

Ministries of Health (MOHs) and non-governmental organizations operating in the field often lack accurate information about settlements at the house- and household level. Detailed maps of where people live and how to reach them are critical for humanitarian aid delivery and infrastructure planning. Organizations increasingly seek this data to run effective operations and facilitate interventions, such as polio vaccination (Price 2015) and malaria prevention campaigns (Price 2017). Figure 1.1 demonstrates the common information gap in a rural region of Burera, Rwanda. Entire villages and major roads, visible in raw satellite images (top box in fig. 1.1), are not represented in the structured map layer (bottom box in fig. 1.1, showing Google Maps as an example). This data gap precludes ways to build digital tools based on geospatial data. One cannot generate travel route descriptions, find homes, or even estimate population distributions without these structured maps. These gaps exist throughout poor areas across the globe, home to millions of people (Nuwer 2014). Many human settlements at the base of the socio-economic pyramid (BOP) are not represented in today's structured map systems. They are digitally invisible.

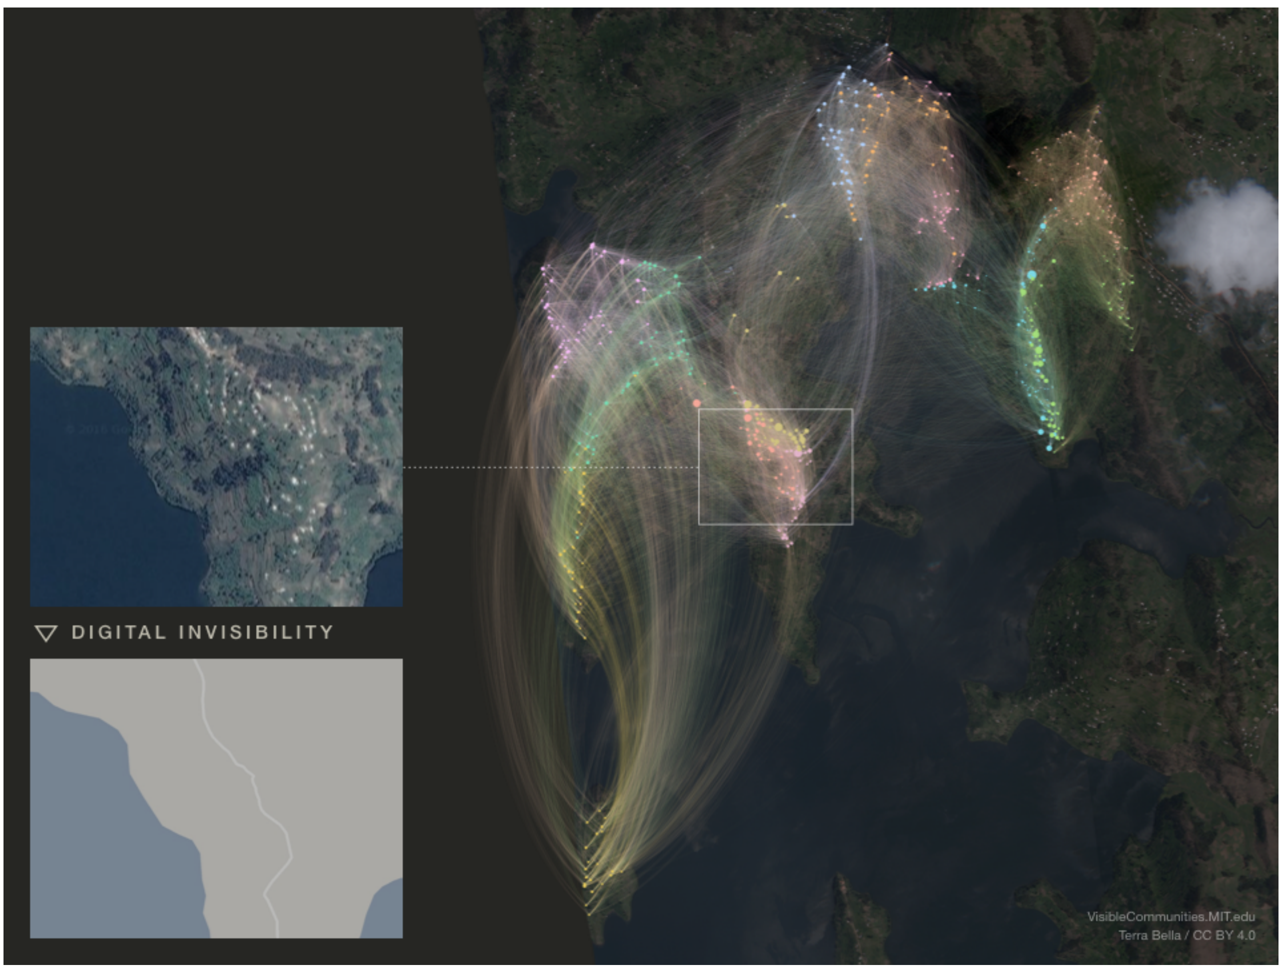

Figure 1.1: Settlements visible in raw satellite images (top box) are not represented in structured map data (bottom box). A socio-spatial map (right) makes such communities visible and expands the definition of being visible.

Digital invisibility is a result of the cost and complexity of existing field-based mapping practices. Even if the tools to create ground-validated maps were available at the BOP, they are not accessible to communities with lower technical literacy. Humanitarian organizations are increasingly relying on satellite imagery of the earth's surface to annotate visible infrastructure. Interpreting remotely sensed images is a practical way of establishing baseline maps that support the operations of such organizations. However, these maps are limited to showing the top-down perspective: they might miss critical ground information, include mistakes, and have biases; they need bottom-up validation.

1.2 The Potential of Human–Machine Collaboration

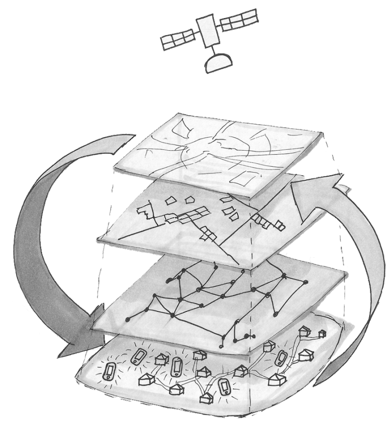

Mapping practices involve human labor and computational power to create detailed maps. Digital tools can assist humans with the complex task of mapping and coordinate the work of large groups of people. Computers can help analyze massive amounts of unprocessed data from satellites and utilize it for mapping. Machine learning can further produce house- and road-level maps using high-resolution satellite images of large regions (further discussed in section 2.1). Automated maps help to estimate where data gaps exist and can serve as scaffolding to create more detailed maps. As these maps only represent a top-down view and are not ground-verified, they are prone to including mistakes and biases. Such maps are limited to the features the algorithms extract from remotely sensed data. Hence, this top-down perspective alone does not reveal the function of a building nor the social characteristics of a community.

Local communities, who are the most knowledgeable about their environment, are best equipped to provide this bottom-up perspective. Tools and programs that enable bottom-up participation can utilize the power of the crowd. The increasing smartphone adoption in low-income economies (see section 2.2) presents an opportunity for large-scale crowdmapping. Low-cost smartphones include capabilities like the Global Positioning System (GPS) and – with appropriate software – offer the possibility to contribute field data to an open map. However, many communities living in poor areas of the world have low literacy rates. Text-reliant interfaces present an obstacle to adopting these technologies. Furthermore, interfaces using standard map representations are inaccessible to novice users who are unfamiliar with classic cartography. We must address these barriers to digital interfaces by designing appropriate interfaces for BOP communities.

Human–machine collaboration can combine the relative powers of each party. Automating the creation of a basic map and crowdsourcing improvements in map quality can quickly create value for local communities and organizations operating in unmapped regions. This two-way integration of top-down satellite data and bottom-up field data creates an interactive, continuously improving spatial map, and can add the social dimension.

1.3 The Role of Community Health Networks

Understanding existing social structures is essential to the successful adoption of new technological systems. We conducted a two-week field observation in Rwanda in September 2016 to learn about the needs and capabilities of communities at the BOP. We visited a rural district where Partners in Health (PIH), a global health organization, supports public health and development activities (Cancedda et al. 2014). PIH operates in three rural districts, supporting three hospitals, 40 health centers, and a network of 7200 village-based community health workers (CHWs). During our field observation, we identified CHWs as ideal ground actors because they are embedded in the villages and trusted by the local communities.

CHWs provide basic health services in villages and are elected from within their communities. The term to refer to this role varies across the world. The World Health Organization (1989, p. 6) defines CHWs as:

> Community health workers should be members of the communities where they work, should be selected by the communities, should be answerable to the communities for their activities, should be supported by the health system but not necessarily a part of its organization, and have shorter training than professional workers.

In Rwanda, elected CHWs are trained by the MOH to deliver basic primary health services, refer sick patients to nearby health facilities, and routinely visit all households to monitor health at the village level. PIH supports this program by employing full-time community health coordinators and creating tools, such as household registry books, reports, and schedules.

PIH is aware of the value of mapping for deploying health services. In 2014, the PIH Research Department trained supervisors of CHWs in one Rwandan district to map village center coordinates using professional GPS devices (Munyaneza et al. 2014). PIH employed a geographical information system (GIS) specialist who used the coordinates to determine the accessibility of health facilities of each village. Based on these results, the MOH constructed three new health centers and dispatched two new ambulance cars within three years of the study (Munyaneza 2016).

1.4 Designing a Social Machine

This thesis proposes a "social machine" design for sustainable and scalable house-level mapping. Hendler and Berners-Lee (2010) introduce a social machine as a system of humans and machines collaborating to accomplish a task that one part could not accomplish alone. Our system's design combines (1) automated satellite image analytics at the house-level, and (2) an intuitive smartphone app for self-mapping spatial and social ground truth. The satellite maps are focused on rural areas in low-income economies, and the app is designed for the communities in these regions. Our approach combines the elements of machine-driven analytics and human-powered annotations into a holistic social machine. Specifically, our proposed design leverages machines to use automatically extracted building footprints to provide a baseline map and enables humans to annotate house-level data in the field. Using low-cost smartphones and an intuitive mobile app accessible to local communities enables self-mapping.

Our approach is designed to be utilized by the local communities being mapped. We focus on end users embedded in these communities to build upon their local knowledge and to provide direct value to the users engaged in self-mapping. The goal is to provide a user experience that aligns institutional interests with grassroots interests, resulting in a self-sustaining system. Specifically, we aim to empower PIH-supported CHWs to create a house- and household-level map of their catchment area by using smartphones. We attempt to design a barrier-free app and expose CHWs to different views of their collected data to encourage greater map literacy. CHWs are credited as authors of their mapped data in the produced maps that are shared with the community. Through design and human engagement, we hope to create the missing bottom-up link.

The top-down perspective can map buildings and roads, and the bottom-up perspective offers a richer view of the ground. For example inhabitants and their relations, the function of a building (e. g. school), its height (e. g. for population estimates), names of points of interest, and places with local significance (e. g. landmarks). We are interested in a map that captures data at the intersection of *social* and *spatial* features to create novel maps, enable improved capabilities for institutions such as global health organizations, and power versatile apps for end users.

Both social and spatial maps can be thought of as networks, or mathematical graphs, with two types of elements: *nodes* and *links*. The foundation of each network is its nodes. In our context, spatial nodes represent homes in villages and social nodes the households living in them. The spatial links are paths connecting the homes, and the social links are the relationships between households. Figure 1.2 summarizes these four map elements: homes, households, paths, and relationships (compare to the generalized socio-spatial map framework of fig. 5.1 discussed in chapter 5). Our social machine is designed to map homes and households, and the thesis further explores paths and relationships through analytics.

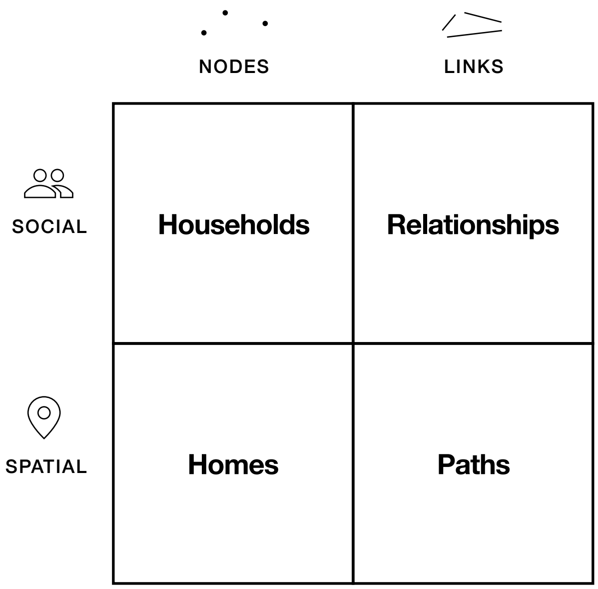

Figure 1.2: Diagram of social and spatial map elements.

This thesis argues that human–machine collaboration empowers digitally invisible communities to make themselves visible. In particular, health workers in rural Rwanda independently used a self-mapping tool we designed to enrich machine-generated maps of their communities. These maps captured not only the spatial dimensions of traditional maps but also included social aspects to create a socio-spatial map that expanded the definition of being visible for these communities.

In chapter 2 we review relevant literature in the fields of mapping methodologies, smartphone adoption in low-income economies, and low-literacy interface design. Chapter 3 describes the design, implementation, and deployment of our system to source novel spatial and social data. Chapter 4 evaluates this collected data, shows ways in which these maps make communities visible, and how the user experience bridges knowledge gaps and opens opportunities to gain new insights for the users. Chapter 5 introduces a generalized framework for socio-spatial maps, discusses their potential, and examines design characteristics of our social machine system. We conclude in chapter 6 with a summary of results and possible future applications.

2
Background and Related Work

We discuss the significance of a self-mapping tool for communities at the BOP and survey the landscape of existing tools. Different mapping methods involve both human effort and computational power for the laborious and complex process of creating detailed maps. We review the utility of satellite imagery for mapping and systems that enable large groups of people to map in the field. While increasing smartphone adoption in low-income economies presents an opportunity for large-scale crowdmapping, barriers such as low literacy remain and must be addressed by designing appropriate solutions.

2.1 Mapping

Detecting Built Structures from Satellite Imagery

Remote sensing is a practical way of acquiring ground surface data. Compared to field-based data collection, remote sensing technologies, such as satellites and aircraft, offer large geographic coverage and repeated measurement at decreasing operational costs. An additional benefit is the ability to collect data from regions that are difficult to access. Given these properties, sectors such as agriculture and land development routinely leverage geospatial data to extract useful information from geospatial data.

With increases in imaging resolution, it is possible to identify individual, human-scale structures anywhere on Earth remotely. This opens new applications for multiple sectors including international development. In a pioneering project in 2005, the American Association for the Advancement of Science analyzed satellite imagery of settlements in Zimbabwe to find government-ordered demolitions of communities that supported the opposition party ("AAAS News and Notes" 2006). The Satellite Sentinel Project prominently used satellite imagery from 2010 to 2015 to document human right abuses in the border regions between Sudan and South Sudan (Harvard

Humanitarian Initiative 2011). In addition to such targeted investigations, high-resolution satellite imagery can also be used to map at the house- and road-level without physically surveying these regions. As discussed in section 1.1, such maps are increasingly created for humanitarian aid delivery.

Greater image availability and higher image resolutions pose the challenge of interpreting massive amounts of new data to create digital maps. There are three approaches currently used to address this challenge:

1. Human annotations: Operators annotate images manually using custom-built software.[1] While some companies charge for this labor-intensive service, the open mapping platform OpenStreetMap and the mapmakers for the navigation app Waze solicit volunteers. Google Maps integrates both approaches by enabling power users to contribute freely to an otherwise professionally created service (Google 2017).

2. Machine predictions: Computers automatically extract structured data from images. This approach is enabled by the rise of machine learning[2] and recent results for building footprint extraction (Yuan 2016) demonstrate the potential to automate the annotation task to a fraction of the time and cost.

3. A combination of the two approaches above results in a semi-automated approach. The company Ecopia, for example, uses computers to first auto-generate maps for large regions, which are then reviewed and improved by numerous human operators.

[1] This process of remote mapping without physically surveying the area is sometimes referred to as "armchair mapping" because it can be accomplished without leaving the chair.

[2] Machine learning methods have become more widely implemented due to advancements in the underlying algorithms and lower cost computing power.

The Missing Maps project[3] proactively adds the vulnerable populations onto a map by obtaining human annotations of satellite imagery during volunteer mapping parties, also known as "mapathons" (Murgia 2016). Such maps were used during the 2014 Ebola outbreak in West Africa to monitor the spread of the virus from hotspots and plan interventions.[4] The Facebook Connectivity Lab uses machine predictions on satellite imagery to create higher resolution population maps by combining the building location data with census data (Gros and Tiecke 2016). These maps are used to direct autonomous airplanes to provide Internet access where people live as further described in section 2.2. The satellite operator DigitalGlobe uses its imagery for mapping projects in both categories: Tomnod is a web-based platform for manual annotations and SpaceNet is a research challenge to develop computer vision algorithms for automated annotations. Examples of mapping tools that leverage these different approaches are listed and compared in table 2.1.

[3] Missing Maps is a collaboration between the American and British Red Crosses, the Humanitarian OpenStreetMap Team, and Médecins Sans Frontières. They primarily contribute to OpenStreetMap, but also developed MapSwipe, a mobile app to crowdsource coarse annotations.

[4] Disease mapping can be traced back to the 1854 cholera outbreak in London's Soho district when physician Snow (1855) drew a map of water wells and deaths and found that all deaths were clustered around a specific pump. This map started modern epidemiology.

Resulting data from any of these approaches reflects a top-down perspective. The data is not ground-validated and by itself limited to the kind of features that can be recognized in satellite imagery. Varshney et al. (2015) demonstrate the potential of using machine predictions to locate households living under thatched roofs for targeted rural development. However, when comparing the predictions with survey data obtained in the field, there was a low correlation between the two. In analyzing the cause, they found that many households lived under iron roofs and constructed thatched-roof structures next to the main house for auxiliary purposes.

Crowd Field Mapping

Creating maps from the ground requires the mapper to be present physically in the area being mapped. Such field surveys enable ground data collection of geographic coordinates and local information. Many aspects of geospatial features, such as their names or functions, can only be mapped from the ground. Field mapping is labor-intensive and can only be accomplished either over an extended period or by utilizing a large number of people. Therefore, it is an expensive kind of mapping if not leveraging the power of a volunteering crowd.

In section 2.1 we describe how crowds increasingly create maps from the top-down perspective. Bottom-up mapping, however, requires physical access to the area and thus crowds are harder to mobilize to cover large regions. Creating general-purpose basemaps (i. e. maps that include the geography and general infrastructure landmarks like buildings and roads) require equipment and skill to accomplish the task in the field. Hence, crowd field mapping typically focuses on creating thematic maps to show a particular theme on an existing basemap.

The activist mapping platform Ushahidi helped produce a crisis map during the 2010 earthquake in Haiti. The project impressively demonstrated the potential of enabling crowds to report the reality on the ground around specific issues during an event. Technology-savvy volunteers on the ground and from remote locations aggregated the reported data into map overlays using OpenStreetMap as a basemap. These efforts closed key information gaps during the first days before humanitarian organizations were operational (Morrow et al. 2011). The Flocktracker platform focuses on coordinating crowds for urban data collection. The MIT Department of Urban Studies and Planning developed this platform to map the semi-formal bus system in Dhaka, Bangladesh by continuously tracking bus movements (Zegras 2013). For disconnected environments, the MapJack[5] mobile app enables communities to create topical maps using low-cost

[5] Not to be confused with the street-level mapping platform similar to Google Street View also called MapJack.

smartphones. Two communities that are frequent contributors to OpenStreetMap used the app in 2015 to map what is locally significant to them (Ntabathia 2015). Ushahidi, Flocktracker, and MapJack rely on existing basemaps such as Google Maps or OpenStreetMap.

Most mobile location data collection apps that can create basemaps such as Esri Collector or Fulcrum fall into the professional GIS category. Data collection efforts using such systems require professional knowledge to setup and utilize in the field. To georeference household surveys, Tata Trusts in India uses the SocialCops app. The additional step to create the mapping layer requires data connectivity and more specialist skills (Garg 2015). A different kind of basemap can be created with Mapillary, an app that crowdsources smartphone photos and processes them with computer vision into a street-level mapping layer for OpenStreetMap. Communities can apply online for a humanitarian mapping kit that provides the required equipment to capture local street-level photos (Uddbäck 2016). Community-based mapping projects like Map Kibera in Nairobi, Kenya and Ramani Huria in Dar es Salaam, Tanzania educate young community members to collect field data with a variety of tools. GIS experts then process that data into accurate digital maps to share with the community.

Participatory mapping practices where lower technical literacy suffices include hands-on mapping techniques such as drawing from memory in sand or sketching on large paper. If basic geographic maps or satellite imagery can be printed on paper, the paper maps can be used by the community to transcribe their local knowledge and fill blank areas. More intuitive, physical three-dimensional models can be created with local materials to map the environment. While local communities can perform such practices without the engagement of experts, the resulting maps are not digital and have lower fidelity (Corbet 2009, p. 45). Ntabathia (2015) argues that only 0.7 percent of all OpenStreetMap contributions come from mobile phones because the digital tools used for participatory mapping were originally designed for the desktop and the web.

Leveraging Local Capacity for Data Collection

Global health organizations are increasingly leveraging local capacity for data gathering for faster response and more comprehensive coverage. The number of CHWs across sub-Saharan Africa and India alone are projected to approach two million over the next decade (Singh and Sachs 2013). Even though the primary role of a CHW is to deliver basic health services, they also collect certain data as described in our field observations in section 3.1. These records are primarily paper-based, and there is an opportunity to increase the efficiency of the reporting process using digital tools. Because of their increasing numbers and intrinsic trust conferred upon them by their community, CHWs are uniquely positioned to collect data beyond their core domain.

In a study led by PIH, Munyaneza et al. (2014) worked with local CHW supervisors to map village centers in Rwanda using GPS devices. A photo of the resulting map is included in appendix E. Since the government had not mapped the locations of villages then, the PIH study provided the best map of that region. The study compares employing GIS experts to map the villages in two districts against training supervisors of local CHWs to map the northern Burera district. The data quality was equivalent, however, leveraging local capacity halved data collection costs.

Several ongoing and emerging projects use smartphones to support mobile health, also known as mHealth. World Vision created five apps for low-resource settings using the CommCare platform by Dimagi. They deployed these apps in ten countries across Africa, India, South and Southeast Asia to aid frontline health workers (Dimagi 2014). Under the Livelihood Empowerment Against Poverty program in Ghana, the app Insyt by Esoko surveyed 150 000 poor households. Three hundred local field agents used the app for data collection (Ntiamoah 2016). A project at the University of California, San Francisco is developing an app for local health workers and active surveillance teams in Swaziland to enter malaria-related data, such as cases of the disease and mosquito breeding sites, to refine risk maps (Kurtzman 2014). A common theme of these projects is that they employ workers who know how to use apps with standard interfaces. Such frontline workers often come from urban areas and travel to remote areas to collect the data. Thus, those projects do not strictly leverage local capacity.

Overview of Mapping Methods

Table 2.1: Overview of different mapping methods with examples spanning commercial services, volunteer projects, and active research.

Satellite imagery			
Human	Machine	Field survey	Satellite & Field
Missing Maps	Ecopia	SocialCops	OpenStreetMap
Tomnod*	SpaceNet*	Map Kibera	Ushahidi
		Flocktracker	Google Maps†
		Mapillary	Waze

* Project by DigitalGlobe
† Includes Local Guides program (previously Google Map Maker)

2.2 *Smartphone Adoption in Low-Income Economies*

While smartphone adoption is rising worldwide, in the context of this thesis, we consider low-income economies in particular. We examine the factors for adoption at the BOP, the timeline, and the potential this technology adoption offers. Because smartphones combine the capabilities of the mobile phone with functionalities that previously required dedicated devices, the change a smartphone can bring about is on a different scale than merely improving an existing technology. A single device uniting the alarm clock, portable audio player, and navigation device, for example, has both cost and complexity advantages. The biggest upside for a smartphone owner, though, is when these functions are not just combined, but made accessible for the first time through the aggregation of functionality. If adopted, capabilities like GPS are – almost as a byproduct – placed in billions of people's hands.

An India-based project by the non-profit organization SocialEmergence reveals several factors that impact smartphone adoption in low-income economies: there is affordability, access to the devices, and willingness to pay. Prices for devices offering the latest technology continue to decline rapidly through competitive market forces. The shift in manufacturing to low-income economies drives the price below $100 per smartphone handset, and even as low as under $50. The supply of low-cost smartphones is further increased by an emerging secondhand market. Device sharing – for example within families – broaden access to acquired phones.

However, the expansion of cell tower infrastructure is expensive and slow. Mobile data connectivity remains unavailable or unaffordable for most BOP communities. While mobile phones bring voice communication to people who never had and never will have a landline phone, smartphones without data and regional content have a weaker value proposition. Sometime in the near future there likely will be a point where the cheapest mobile phone will be a smartphone as we define it today, but with the simple and robust budget phones that are available, it is unclear when that will be.

Technology companies are exploring unconventional ways to accelerate the coverage of large rural areas in low-income economies because mobile connectivity is key for the adoption of smartphones and their services. Google Loon is a network of high-altitude balloons that send cellular data connectivity to the ground. Similarly, Facebook's Connectivity Lab is testing Aquila, a network of high-altitude solar-powered autonomous airplanes to address this challenge. At even higher altitudes, companies are investigating the use of satellites. Such efforts are still nascent. On the ground, the increasing density of smartphones and Wi-Fi routers have the potential to transmit data through wireless mesh networks. The seminal book on the topic Wireless Networking in the Developing World (The WNDW Authors 2007) describes building such networks as technically feasible, but continuing operations a more difficult business challenge. In multiple case studies, the authors discuss the importance of local community involvement.

While smartphone subscriptions are increasing rapidly, most users at the BOP have yet to be reached.[6] Africa is the second largest region behind Asia Pacific of mobile phone subscribers, yet it is also the one with the lowest adoption (GSMA 2016, p. 8). From the supplier side, the number of smartphones surpassed the number of traditional mobile phones, and by 2020 there will be a projected 726 million smartphone users in Africa (GSMA 2016, p. 13). However, the adoption will be throttled – especially at the BOP – by connectivity issues, illiteracy, and therefore lack a compelling use case for these communities. The barriers to adoption will shift from device availability and affordability to a clear BOP end user value proposition. Intentionally designed software has the potential to address these remaining barriers.

[6] Over 80 percent of African countries are classified as low- or lower-middle-income and 27 out of worldwide 31 low-income countries are located in Africa (World Bank 2017).

2.3 *Low-Literacy Interface Design*

Illiteracy is generally understood as the inability to read and write text. The United Nations Educational, Scientific, and Cultural Organization defines functional literacy as the ability to read, write, and calculate sufficiently to function in one's local community. According to the UNESCO Institute for Statistics (2016), 758 million adults worldwide were functionally illiterate in 2016. In addition to text literacy, there are other literacies we consider in this thesis such as digital literacy and map literacy. In this section, we explore in particular the use of digital interfaces for text illiterate users.[7]

Research on the use of mobile phones in countries with low literacy rates by Chipchase (2005) from Nokia concludes that phones for illiterate users should not be noticeably different from established products on the market to avoid the social stigma associated with illiteracy. Purely icon-based user interfaces (UIs), furthermore, are not regarded as the solution.[8] Designers should draw from visual cues[9] that the users are familiar with through local knowledge and experience. Even with the absence of familiar cues, the steps to carry out most tasks can be successfully memorized. However, memorization does not mean understanding, which is usually required to solve new problems. Chipchase argues that the users' willingness to explore the UI is correlated with perceived risk of consequences. For example,

[7] CHWs in Rwanda that participated in our study were literate in the local language, but we assume low literacy for a scalable and inclusive solution. Programs that require health workers to be locally selected but lack strict education requirements may lead to decreasing rates in CHW literacy.

[8] Wiedenbeck (1999) shows that the combination of icons and labels in UI design yields the best results for learning and retention. Although icons-only performed poorer in the early stage of learning than text-only, users favor UIs with icons.

[9] Bank notes are a good example of offering multiple cues to understand their value.

if there are multiple options and one option might delete the user's data, users are not likely to engage with the interface.

A large body of published work on text-free interfaces from Microsoft Research India demonstrates the effectiveness of hand-drawn, semi-abstracted cartoons with voice annotations for illiterate users (Medhi, Sagar, and Toyama 2006; Medhi Thies 2015). Although this technique was designed for novice and illiterate users, it works most effectively when human mediators guide the users. The Microsoft Research India team also recommends interfaces that do not scroll and limiting text labels to numbers only. These recommendations run counter to the other reviewed literature.

In a study of illiterate users, Knoche and Huang (2012) show that the use of some text in the interface helps illiterate users to filter and group data like a phone's address book. The authors' research group at École Polytechnique Fédérale de Lausanne also created the mobile app EasySMS in 2011 that enables illiterate users to send and receive text messages. The user can interactively hear messages through a word-by-word UI inspired by Karaoke. Audible icons of common words and recomposing of words from received messages enable the user to reply to messages. EasySMS is a notable example of an illiteracy-proof UI that encourages content creation.

In summary, the reviewed literature shows the benefit of less text-reliant UIs for novice and illiterate users. Audio-visual cues can support the accessibility of UIs.[10] These cues should be used to aid the user, and the interface should not be reduced to just those cues for usability and social reasons. Simple and intuitive interfaces are preferred by illiterate users, which is also true for literate users. Good practices include:

- Focused feature set with one user action per feature[11]
- Minimal hierarchical navigation
- Consistent visual language

[10] Recent research opens the opportunity to use speech interfaces for low-resource languages. Babel is a project that aims to broaden the set of languages for speech recognition. Google Research is crowdsourcing training data to achieve the same for text-to-speech (Ha 2015).

[11] For tasks that cannot be accomplished with a single action, a path to delegating subtasks should be available.

2.4 *Significance of a Community Self-Mapping Tool*

Digital invisibility is a problem for communities at the BOP and, as discussed in chapter 1, the data gap will likely persist due to the challenges for governments and corporations to map. To close this data gap, it is important to empower communities to make themselves visible. Humanitarian mapping efforts have the tendency to concentrate on crisis response instead of sustained rural mapping. In addition to practical reasons for a community to be visible, there is a natural and universal desire to be recognized. When a community has the tools to self-map, it not only enables but also empowers the community to decide where and how they are visible. While various mapping practices summarized in table 2.1 address aspects of closing this data gap, available solutions to create basic house-and road-level maps are not ground-validated, do not use digital tools and are not accessible to local communities.

Solutions for social self-mapping – a different kind of self-mapping – remain largely unexplored. By social mapping, we mean digitally establishing social links. While Internet users voluntarily socially self-map using services such as Facebook, Twitter, and LinkedIn, much of the underlying data is inaccessible to users and researchers. The MIT Media Lab community created a small social network from the ground up using the self-mapping web tool proposed by Saveski et al. (2016). In the context of international development, studies by The Abdul Latif Jameel Poverty Action Lab at MIT have shown the potential of using fine-resolution social data of BOP communities to influence positive behavioral change and to fight poverty. For example, Banerjee et al. (2013) conduct a rural village-level randomized controlled trial and demonstrate the relative impact on an entire community depending on which members receive microloans. Similarly, Kim et al. (2015) exploit properties of human social networks[12] to target central members of rural villages and demonstrate a significant increase in adoption of health interventions for the entire community. Organizations studying such data must either administer an ad hoc survey with trained labor in a traditional way or hire a company such as Premise Data that provides surveying services. However, this way of surveying communities is uneconomical and does not empower the communities directly.

[12] Namely, they target nominated friends of randomly selected villagers to exploit the so-called friendship paradox; colloquially, "your friends have more friends than you do."

The participatory mapping study by PIH described in section 2.1 showed how local capacity was used effectively to map. The key differences between this study and the tool we propose include:

- mapping at a house-level resolution (compared to village-level with a single point per village),
- using localized software designed for low-literacy running on commodity smartphones (compared to professional GPS devices in English),
- giving the tool to CHWs directly (compared to their supervisors),
- aligning the activity with the users' natural workflow and motivations (compared to training them for a specific task),
- collecting spatial and social data (compared to exclusively spatial data).

The near-term global availability of smartphones and high-resolution satellite imagery present an opportunity to close this data gap economically and at scale. For the human side to scale, a solution must be designed at the local community level. A local partner is essential to connect technology with particular communities. Community self-mapping offers the opportunity to not only map spatial infrastructure but also to capture social relations that are invisible from a top-down perspective. Human ground annotations are richer and can be credited to people living in the mapped area, creating a desirable bottom-up dynamic. Thus, we set out to research the social and technological aspects required for a human–machine collaborative system that bridges the existing gap and yields a useful map for communities, organizations, and the broader world.

3
Methods and Implementation

We completed a two-week pilot of our house- and household-level mapping system working with PIH and twelve local CHWs in early 2017 in rural Rwanda. The health workers mastered the use of our mobile mapping app Yego with minimal training, integrated it into their daily routine, and mapped their entire community of over three thousand people in short order. The voluntary, bottom-up community mapping was aided by machine-generated predictions using custom acquired satellite imagery from SkySats. We designed the app UI based on background research and insights from prototyping during a two-week field study in mid-2016. We aggregated the data into a novel map of that region and shared it with PIH and the local community. We use these results together with CHW interviews to evaluate our system.

3.1 Field Observation in Rwanda

In September 2016, we conducted a two-week field observation in Rwanda to understand the needs of communities at the BOP and organizations supporting them. We visited the rural Burera district where PIH supported public health and development activities as introduced in section 1.3. Our goal was to identify a stakeholder on the ground with a functional need for maps, learn more about them, and prototype a potential solution. Furthermore, we sought to establish a relationship with an organization to run a pilot with us jointly. Rwanda is a low-income economy, according to the World Bank (2017) index, and 83 percent of the population live in rural regions (National Institute of Statistics of Rwanda 2012, p. 9). These regions were largely unmapped: from an outside perspective, the location – and even existence – of most buildings and many roads were not known.

Partners in Health's Operation

PIH built and operated three hospitals in three districts in Rwanda. About a dozen smaller public health centers surrounded the hospitals. Rwanda's 30 districts further subdivided into 416 sectors and 2148 cells. A cell typically contained 5 to 15 villages.[1] Community-based government programs were organized locally at the cell-level. PIH supported community health programs by employing health center staff and three full-time community health coordinators country-wide. These coordinators facilitated training for CHWs and managed cell-level CHW supervisors. PIH created tools for the CHWs such as household registry books, reports, and schedules.

[1] The smallest administrative entity in Rwanda is called *umudugudu* (pl.: *imidugudu*) and is an agglomeration of houses comparable to a village in other countries (Ministry of Local Government 2011).

The Butaro Hospital in the northern district of Burera was surrounded by 573 villages, each typically consisting of 50 to 150 households. This mountainous northern region contained rural settlements and was difficult to navigate via motorized vehicle. As described in section 2.1, the PIH research department completed a project that trained CHW supervisors to obtain the center location of all villages with a GPS device. PIH implemented this project to link insights from GIS to health outcomes. Other practical needs for detailed maps of the regions included planning and resource allocation of efficient healthcare service delivery.

Ethnographic Study

In our field observations, we immersed ourselves in the communities with the objective of understanding the culture and systems relevant to designing an appropriate solution. During our work, we were also conscious of the 1994 genocide that formed present-day Rwanda. At the time of the study, the dominant social norm was that all Rwandan are one national group and the census no longer tracked ethnicity. The focus of this section is on the social and technological findings relevant to our work in the context of this thesis.

Most of the population in Burera were subsistence farmers. A minority of those farmers owned a few cows, goats, or sheep. Occupations away from home, for example in the road construction business, were rare and part-time. Families farmed on fields around their homes, resulting in scattered settlements. The local house density was sparse, but the population density over the entire region was high. Other than topography and homes, there were few landmarks. The farmers usually stayed close to their homes and fields. Family members – mainly children – fetched water daily at central water sources, often an hour-long walk away. On average, once a week the farmers commuted to a regional market and occasionally to the nearest health center. Electricity only reached the commercial centers – an aggregation of a few shops with basic items. Shops offered electricity stations for people to charge mobile phones. People carried feature phones on the street, and such phones were usually the first electronic device a family owned. Mobile phone companies were often the only commercial offering alongside roads. These booths sold or rented low-end smartphones as well. The cell phone coverage was decent but cellular data connectivity mostly unavailable. Photos of the region, a village scene, a commercial center, and a phone stand can be seen in appendix E.

As described in section 1.3, CHWs are elected from within their communities.[2] They were subsistence farmers like the vast majority of people in the village. We worked closely with CHWs because they were members of the community who we could be introduced to through their tie to PIH. We met with six CHWs at local health centers or their homes to learn about their specific tasks, what their interactions with the community entailed, and how they coordinated their work. We also sought to understand their daily routines and motivations.

The CHWs we spoke to had on average served ten years in their roles with low turnover rates. Because they were all long-time members of these communities, wayfinding was not an issue for them. They were familiar with their surroundings and oriented themselves from memory. Verbal directions were often difficult to obtain because of the lack of landmarks and the custom of showing someone the way instead. When asked to sketch a representation of their village (see examples in appendix B), about half of the CHWs drew from their perspective standing at ground level as opposed to conventional top-down cartographic representations. When showing them the map PIH created of their region, they recognized neighboring villages based on the printed names. The local administrative boundaries were well-known because they were defined according to local geographic features.[3] These local features helped in defining a CHW's household catchment area. A strict inner-village division of households was coordinated individually by the CHWs to ensure no household is missed.

Filling out reports constituted one of the key criteria to be nominated as CHW. All CHWs attended a few years of primary school and could read and write in the local language Kinyarwanda.[4] They were all respected and trusted by the communities they served. Their motivations to be a CHW were to learn new skills and help the community. They received a monthly financial incentive of 5000 Rwandan Francs[5] from the MOH, which according to all CHWs was negligibly small. However, their role enabled them to work with leaders from the local governments. Some CHWs saw an opportunity in receiving

[2] There are usually three CHWs per village: one female and one male called *binôme* (from French, "a pair") responsible for general community health and a female maternal health worker called *ASM* (from French, "Animatrice de Santé Maternelle"). Populous villages have three *binômes*.

[3] In 2006, the government determined new village boundaries in a bottom-up process: cell authorities took natural features (e. g. valleys, streams, forests) into account when recommending boundaries to their sectors (Ministry of Local Government 2006).

[4] In addition to the local language Kinyarwanda, people who went to school before 1994 tend to speak French. Since then, English is increasingly taught. CHWs spoke neither at a conversational level.

[5] At the time of the study, 5000 Rwandan Francs were approximately $6.

the training to advance their knowledge to better care for their families. However, the main motivation across all observed CHWs was that they were happy to care for members of their community.

CHWs carried out specific health activities such as monitoring the nourishment of children under the age of five, malaria prevention, and accompaniment of HIV/AIDS patients. Additionally, CHWs conducted general health check-ins with families: twice a month, they visited all households in their catchment. Upon completion of such a regular check-in, they recorded the visit in a book by writing the date next to the head of household's name. Each household in a CHW's catchment was listed in that book (fig. 3.1, top left). To organize the households, the CHW assigned a number to each starting with "1." Deciding on the numbering was up to the CHW and often arbitrary. One pattern was that the CHWs often started by assigning number "1" to their household. Catchments were fixed, and CHWs knew the names of all heads of households.

The CHWs visited homes on their schedule. One of the main challenges for CHWs was the time it took to coordinate meeting families in their homes. If no one were present, CHWs would need to return later. CHWs' days primarily constituted of working in the fields and caring for their families. CHWs also reported that their community health duties were increasing, yet there was no increased pay. If they had received a higher pay, they would have used the money to hire someone to work on their fields and spend more time on the community health activities.

They all received a feature phone through a government program and were adept at using it. The phone was used for voice-based coordination, and text message-based health reporting using a system called RapidSMS.[6] The phone's menus were in French because a version translated to Kinyarwanda did not exist. In addition to the phone, some families owned a battery-powered[7] flashlight or radio. Some have heard of smartphones, but none have used them. Because smartphones were not part of that region's culture, properly handling the device or turning it on posed challenges. When we demonstrated smartphone use, the learning and retention rate was comparable to typical novice users.

We synthesized the observations and mapped out all key stakeholders[8] with their expertise, resources, potential interest, and the dynamics between the different parties. This stakeholder analysis helped us in facilitating an on-site design workshop with PIH staff and CHWs.

[6] RapidSMS was developed by the United Nations Children's Fund Innovation program, and its successor platform that allows for customization is called RapidPro.

[7] We found early-generation solar cells intended to provide ceiling night-lighting and charging devices on a few roofs per village, but all of them were out of order.

[8] Key stakeholders of our analysis: household, CHW, CHW supervisor, PIH CHW coordinator, PIH research department, PIH leadership, health center staff, MOH, MOH Ethics Committee, local government, local vendor, phone supplier, mapping service, satellite imagery vendor, social impact fund

Design Sprint

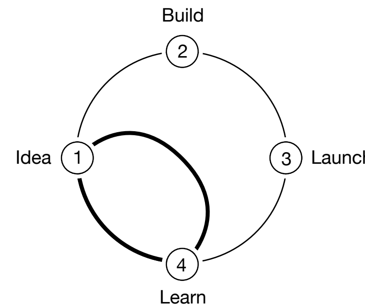

To translate observations rapidly into insights on what solution could work, we planned a design sprint workshop at the PIH headquarters with full-time PIH staff from the community health program. A design sprint is a structured process to directly move from idea to learning, before building a product based on unvalidated assumptions (Knapp, Zeratsky, and Kowitz 2016). The learnings are documented from user reactions to a realistic prototype. GV, the venture arm of Google, who developed the original process, advised us to adapt the typically week-long design sprint process to a three-day time frame and the context of international development.

To cover all roles, we had recruited the sprint team of seven people one month before the workshop. We brought together the PIH staff who has the biggest exposure to the problem of data poverty, staff with the most frequent interactions with the community health program, and staff who work with the involved data and technology. We defined our challenge and long-term goal as:

Empower rural communities to put themselves on the map.

We set out to design a solution to be used directly by members of rural communities, but before the field observation study, it was not clear how feasible that is. We set an initial focus on learning more about the community health system and found that CHWs indeed were ideal actors on the ground. They are regular residents in poor rural villages, present in all communities, yet have a weak institutional link to our partner organization. During our interactions with CHWs before the sprint, we have seen that they were motivated and willing to use new technology. Working towards the long-term goal, we had multiple open questions such as, "Can we address a need of CHWs with a map?" and "How can a mobile app give insight on CHW workload (space and time)?" For our sprint, we chose to focus on the question:

Can we design a mobile app that CHWs are able to use?

Over the course of three days, we mapped out the challenge we set forth, sketched solutions, picked the best ones through a structured decision-making process, and built a realistic prototype to test in the field. We set up a second visit to the rural region that the CHWs are based in to test the prototype. We conducted one-on-one user tests with CHWs in a small usability study to see what works and what does not work. This led to clear actions for the design and development phase back at the MIT Media Lab.

When building a realistic prototype of a complex system in less than one day, it is important to decide on the right tools. The de-

Figure 3.1: Impressions from the field observation: a CHW's list with numbered households (top left), a CHW's first interaction with a smartphone (top right), and PIH staff participating in a design sprint (bottom).

sign sprint process suggests different options for different scenarios such as when designing software, a service, and an object (Knapp, Zeratsky, and Kowitz 2016, p. 186). The tools we chose to prototype our software system with social components were the presentation software Keynote and personal acting. An interactive full-screen Keynote presentation on an iPhone could simulate a realistic app and role play at the CHWs home could simulate a home visit. We had the CHW follow a step-by-step guide simulating onboarding and data entry during a general check-in. We tested whether the CHW could tap through the experience, follow instructions on the screen, and translate them into the desired real world activity with the actors.

Key insights from the ethnographic study, interviews, and design sprint:

- CHWs are trusted by their local communities, who are bound together by strong relationships.
- In addition to case-specific home visits, CHWs routinely visit each household in their catchment every two weeks using a paper list that enumerates all households.
- The main challenge for CHWs is to balance the community health activities with their regular work as subsistence farmers.
- Catchment areas of CHWs are unknown to PIH and would be valuable to know (optimize coverage, detecting anomalies, personalize financial incentive).
- A house- and household-level map has high value (resource allocation and planning, GIS research, measuring impact and presenting to the MOH).
- The northern district of Burera contains rural communities that are accessible from the Butaro Hospital.
- First-time users of smartphones are curious and motivated to learn (sense of delight).
- Tapping is the only discoverable gesture (scrolling, panning, and pinching are learned when demonstrated).
- A video onboarding tutorial showing a person that gives virtual instructions is not intuitive.
- The collection of health-related data should be planned for after the pilot phase focused on mapping (approval process).

In summary, after our field observation, we understood the local needs and had a plan for a pilot with a ground partner. We selected a well-located[9] cell in the Burera district, where PIH would jointly run a pilot when the system was ready. We identified the opportunity to build a mobile app for CHWs that is simple to use and captures geographic data during their regular home visits, avoiding increased demands of the CHWs. The CHWs expressed interest in trying out new tools for their activities, and the resulting data is directly useful for PIH to improve their community support programs.

[9] Our criteria were (1) rural, (2) accessible within a day from the Butaro Hospital.

3.2 *Satellite Data Acquisition*

Images of the earth's surface from above offer a great basis for starting new maps as described in section 2.1. We use such imagery to accomplish two separate tasks: (1) automatically predicting building locations at scale, (2) displaying a visual base layer for the mapping app. The satellite images seen in popular mapping services are commonly licensed from imaging satellite operating companies. These images are highly processed to visually look appealing and are only available for display within those services. To analyze an entire geographical region, and to bundle imagery of the same region with an app, we require access to raw images of our region of interest.

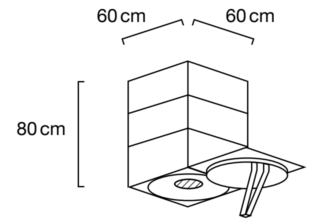

When licensing satellite imagery, two concerns are their recency and cost. Collecting new high-resolution images on the mere city-scale costs in the $1000 range.[10] We purchased archival data from 2014 to run initial tests. The acquired imagery was the same that could be seen on Google Earth and was taken by one of the two Airbus Pléiades satellites. Promising results from our preliminary analytics on the raw imagery data, and the fact that ground reality continues changing led us to look for opportunities to acquire more recent imagery and at regional-scale.

Recent innovations in lower-cost microsatellites[11] are making satellite data more accessible. One such constellation is the currently seven SkySat satellites[12] operated by Terra Bella (formerly Skybox). SkySats reach the high spatial resolution of traditional satellites, and because a constellation contains multiple satellites, collectively offers a high temporal resolution of imagery from daily visits to a site. The capability to image the entire planet every day offers huge potential for commercial and humanitarian use of satellite data.[13]

To access such imagery, we partnered with Google, who acquired Skybox in 2014. Under the Google Earth Outreach program – and recently through Skybox for Good (Mann 2014) – Google contributed fresh satellite imagery to our project.[14] During the last four months of 2016, SkySat satellites captured images of our region of interest around the PIH Butaro Hospital. The clear image we use (right in fig. 3.2) extends from the hospital grounds to the entire largely unmapped surrounding rural region.

When comparing the previously purchased archival patch (mid-2014) with the same subregion in the freshly acquired image (late 2016), we found that the villages evolved significantly in the two and a half years that have passed. Although built structures generally remain relatively constant over time, roof materials change rapidly in developing countries, as homeowners convert thatched roofs[15] into more watertight metal and plastic roofs. A reasonable guideline

[10] The satellite image vendor LAND INFO (2017) quotes prices for new tasking high-resolution imagery in the $30 range per square kilometer (or $15 for archival imagery). For example, the area of San Francisco is $\sim 11 \times 11$ km or 120 km^2 resulting in a cost of $3600 ($1800 respectively).

[11] By engineering high-performance commercial off-the-shelf parts – such as the utilized telescope and imaging sensors – into small cubes, more satellites can be launched quicker. Post-processing the data on the ground using software results in state of the art imaging performance.

[12] SkySat-1, the first of a planned constellation of 24, was launched to orbit in 2013 together with other payload aboard a Russian Dnepr rocket (Terra Bella 2013).

[13] During writing of this thesis, Planet Labs acquired Terra Bella with its SkySats to complement Planet Labs' existing constellation (Marshall 2017) and simultaneously launched additional 88 of their Dove satellites to orbit (Schingler 2017). This makes imaging the entire planet freshly every day possible for the first time.

[14] Planet Labs as well as the DigitalGlobe Foundation maintain similar data grant programs to assist the academic and humanitarian community.

[15] Most satellites infrequently pass over countries close to the equator because they were launched from countries in the northern hemisphere and hence are on a polar orbit.

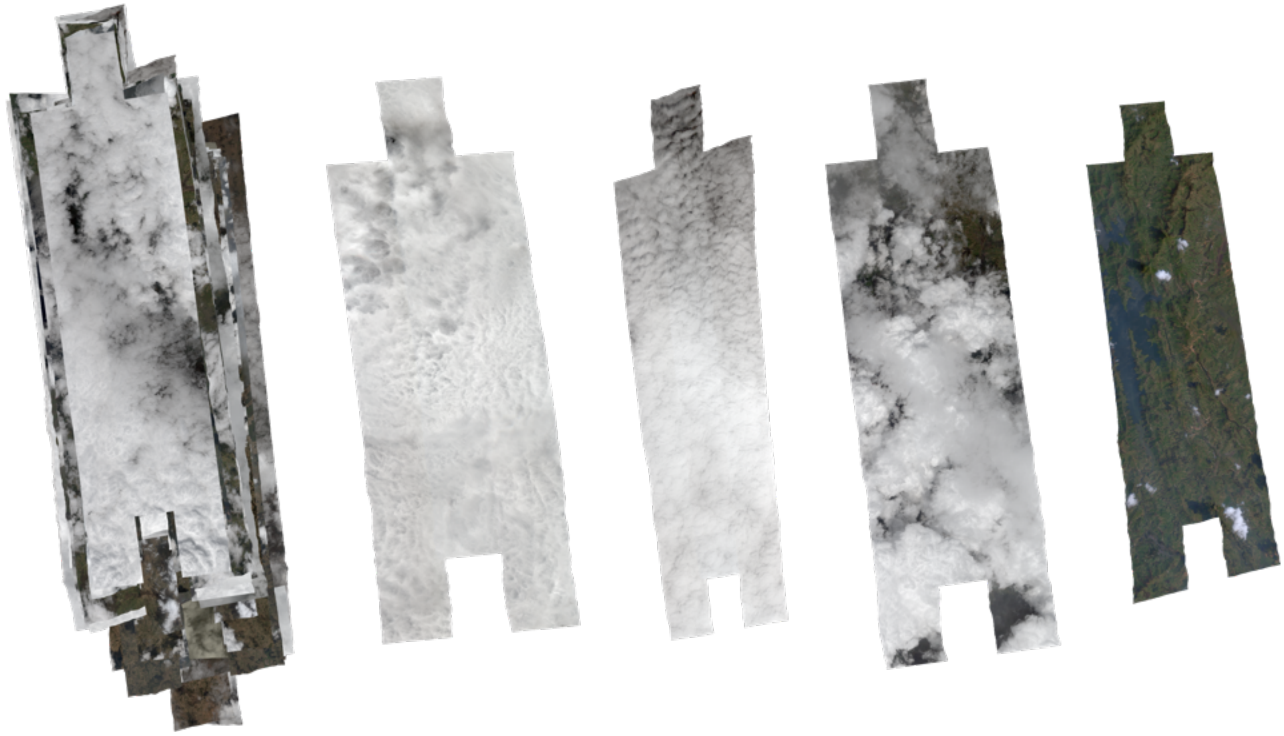

Figure 3.2: SkySats captured 39 satellite images (left), including many with cloud coverage (middle three), until SkySat-3 got a clear view onto Burera, Rwanda in the morning of December 21, 2016.

is to work with data no older than approximately 90 days for built infrastructure assessments. This ensures that we start the ground modeling and later ground validating process with an up-to-date dataset of what is actually on the ground in present-day reality.

It is noticeable how many of the acquired images have strong cloud coverage. As users of mapping services, we are used to seeing the earth's surface clearly. However, when we look up at the sky it is usually at least partly cloudy.[16] It was not until 2013 when Hancher (2013) and a team at Google computed the first global cloud- and cloud shadow-free image derived from Landsat imagery. It took analyzing sometimes dozens of images of a single spot in the world to eliminate clouds and other atmospheric effects. For our regional high-resolution image, we waited for a clear day.

[16] Nearly 60 percent of the total land area is covered by clouds at any given time (NASA 2017).

The image is a systematic mosaic of photos taken from low Earth orbit.[17] The image is generated using "super-resolution multi-image mosaic fusion" from three 2560×1080 pixel imaging sensors mounted side-by-side in the spacecraft (Dirk Robinson 2014). The sensors are offset vertically to make them fit in the focal plane, resulting in a shape with a peg at the top and a slot at the bottom as seen in fig. 3.2. Each of the frames along the three stripes has a ground footprint of approximately 2.6×1.1 km, covering a swath width of 8 km

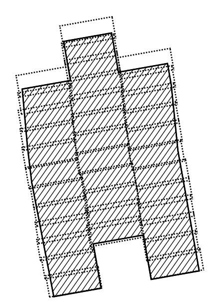

[17] It is like taking a photo of San Francisco from ~ 500 km-away Los Angeles while moving at 7.5 km/ sec and being able to see objects the size of a bicycle.

at nadir.[18] It took 30 seconds of flight to capture our 25 km-long image strip resulting in an area cover of 200 km^2.

In addition to capturing the human visible spectral bands red, green, and blue, to reproduce natural "photographic" color, the multispectral sensor captures a near-infrared and a panchromatic band. The panchromatic band is sensitive to all spectra from blue to near-infrared. The sensitivity at which the sensor captures the electromagnetic energy in these five bands is the radiometric resolution. It is expressed in the number of grayscale levels that can be represented by any of these bands. Typical digital images store the level for each pixel as 8-bit number (0–255), also known as bit depth. In comparison, the raw data in our image is a 12-bit number (0–4095), resulting in a 16 times more nuanced image.

Together with the discussed *temporal*, *spectral*, and *radiometric resolution*, the *spatial resolution* completes the common types of resolutions when considering satellite imagery for remote sensing applications. The spatial resolution describes how small of an object can be recognized and is the core way to benchmark satellite imagery. It is commonly referred to as ground sample distance and defined as how big of a ground surface area maps to a single pixel in the image, expressed in meters per pixel (m/px). The U.S. Geological Survey (2016) Landsat satellites, for example, have been imaging the earth in the 60 m/px (Landsat 1, 1972) to 15 m/px (Landsat 8, 2013) range. Such a ground sample distance shows shapes of landmasses and rough change of their surface.[19] Our image has a resolution of ~ 0.8 m/px[20] at nadir and enables us to recognize all human-built features in the entire region (fig. 3.3, right).

With dimensions of 15 360 × 38 400 pixel, our image is a 590 megapixel photo. Some of our captured images are more than one gigapixel – that is a single image composed of one billion (10^9) pixels. The file size of

[18] Nadir is the point on Earth straight below the spacecraft, resulting in the shortest distance between sensor and imaged surface on Earth.

[19] Today's highest resolution imagery that is commercial availableble at up to ~ 0.3 m/px opens many applications as most daily changing human-scale activities are in this range. Higher resolution than that, for example showing individual people, is aerial photography captured from airplanes and has limited coverage.

[20] In a process called pansharpening the red, green, and blue bands' resolutions are increased from the captured ~ 2 m/px to match the panchromatic band's resolution of ~ 0.8 m/px. To simulate looking straight down from an infinite distance, the image is "orthorectified" by removing perspective and accounting for the terrain.

Figure 3.3: Clear view of a shoreline in Rwanda, Butaro shown at increasingly higher resolution, until small huts and narrow, unpaved footpaths become visible.

Location	Burera, Rwanda (region around Butaro Hospital at 1°24′36″N, 29°50′23″E)
Coverage	$\sim 8 \times 25$ km (200 km^2)
Satellite	SkySat-3 operated by Terra Bella
Date and time	December 21, 2016 at 09:46:16 (local Central Africa Time) during 30 seconds
Spatial resolution	~ 0.8 m/px
Bands	Blue (450–515 nm), Green (515–595 nm), Red (605–695 nm), Near-Infrared (740–900 nm), and Panchromatic (450–900 nm)
Bit depth	12-bit (0–4095) stored as 16-bit (0–65 536)
Image dimensions	$15\,360 \times 38\,400$ px (~ 590 megapixels)
File size and format	~ 5.9 GB ($15\,360 \times 38\,400$ px $\times$ 5-band $\times$ 16-bit) GeoTIFF

Table 3.1: Technical specifications of the satellite image we use for building location analytics and as visual base layer in the mobile mapping app.

our image is ~ 5.9 GB and it is impossible to open it with standard software. Working with satellite data and files this size requires special processes. To create images for display from the raw data, and to extract information from it, we must open and process the image programmatically.

3.3 *Building Location Prediction*

Using machine analytics to extract information from satellite images is a promising technique, as the latest advancements described in section 2.1 show. We use state of the art machine learning techniques on our custom-acquired raw image to automatically predict building locations at the scale of the entire region of interest. The basic idea is to annotate a few buildings and then let machine learning predict the location of all other buildings in the image.

For all image analytics, we use Google Earth Engine, a planetary-scale geospatial analysis platform (Google Earth Engine Team 2015). Earth Engine offers analytics capabilities that can be invoked programmatically and a web interface for visual inspection of the data.[21] The complex operations are then executed on Google's cloud computing clusters. Additionally, the SkySat data can be operated on directly in the cloud without having to download the huge files.

[21] A comparable geospatial big data platform is DigitalGlobe's GBDX.

The raw satellite data has a higher machine interpretive capacity compared to images processed and optimized for the human eye. By letting the algorithms learn from all spectral bands[22] – even ones that are not perceived by the human eye, such as short-wave

[22] Some satellites capture even more spectral bands. For example, WorldView-3 data contains 29 bands including a dedicated cloud detection band (DigitalGlobe 2014).

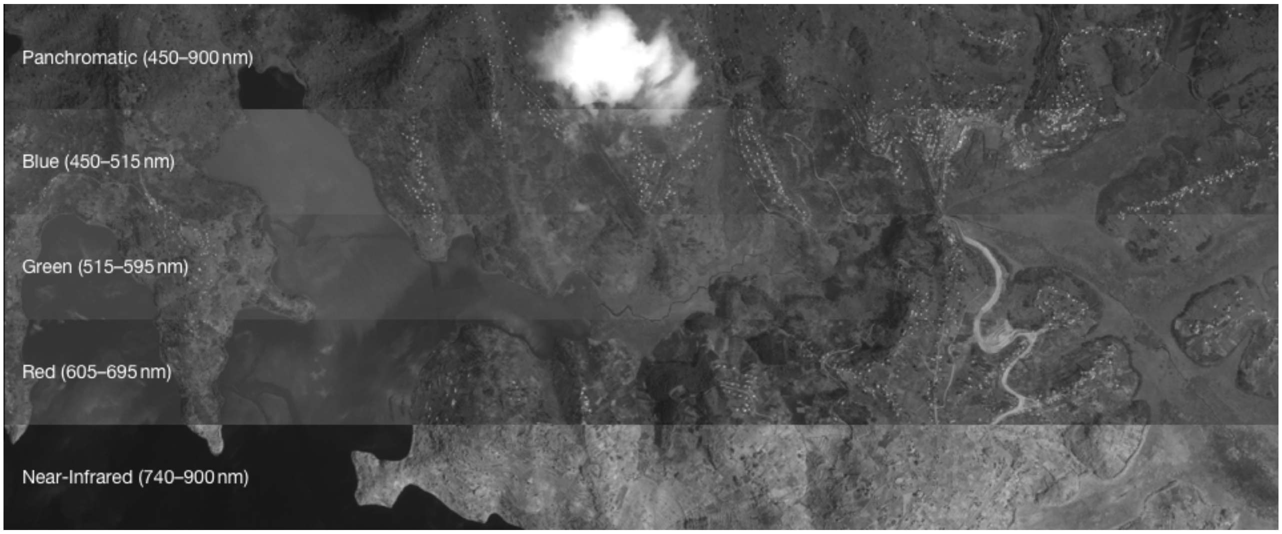

Figure 3.4: All five spectral bands of our image visualized in grayscale. Note water showing up at different intensities (left), cloud and settlements reflecting differently (middle), and the visibility of the road changing across wavelengths (right).

infrared – classification results can be more accurate. The buildings we want to identify have roofs with a distinct spectral signature in multiple bands and can be classified relatively easily.[23] At this time, almost all residential houses are covered by a uniform iron sheet roof.[24] A field photo of a typical home can be seen in appendix E.

First, we manually annotate a few dozen building roofs by placing single-point markers on them at the highest zoom level. We then train a machine learning model with the labeled image area. With that model, we predict building locations on the full dataset (fig. 3.5). Finally, we use image processing techniques to clean the results and turn them into exportable vector data to use in the mobile app.

To decide what technique to choose for our predictive analytics from the library of machine learning algorithms, we considered the kind and number of training data we could collect, and the desired accuracy of the result. Machine learning can broadly be divided into three categories: (1) supervised learning, making predictions based on training examples, (2) unsupervised learning, finding patterns without requiring any labeled examples, (3) reinforcement learning, choosing an action in response to some input and tweaking future actions based on how well the action was chosen. Most problems – including finding all buildings similar to a few collected examples – fall into category (1), supervised learning.

Our goal is to categorize the entire image into two categories "building" or "not building" and hence our problem is best described as a two-class classification. For its accuracy and fast training time, we chose a two-class decision forest classifier (Azure Machine Learning Team 2015). Specifically, we use the random forest classifier

[23] To detect other features (e. g. roads, and buildings that visually appear differently), either multiple training sets and feature classes could be created, or Convolutional Neural Networks could be used to segment an image into different classes automatically.

[24] In recent years, most traditional round-shaped grass-thatched houses *nyakatsi* were replaced with more stable structures covered by an iron sheet. This community development program is called "Bye Bye Nyakatsi" and is part of Rwanda's Vision 2020 campaign (Ministry of Local Government 2010).

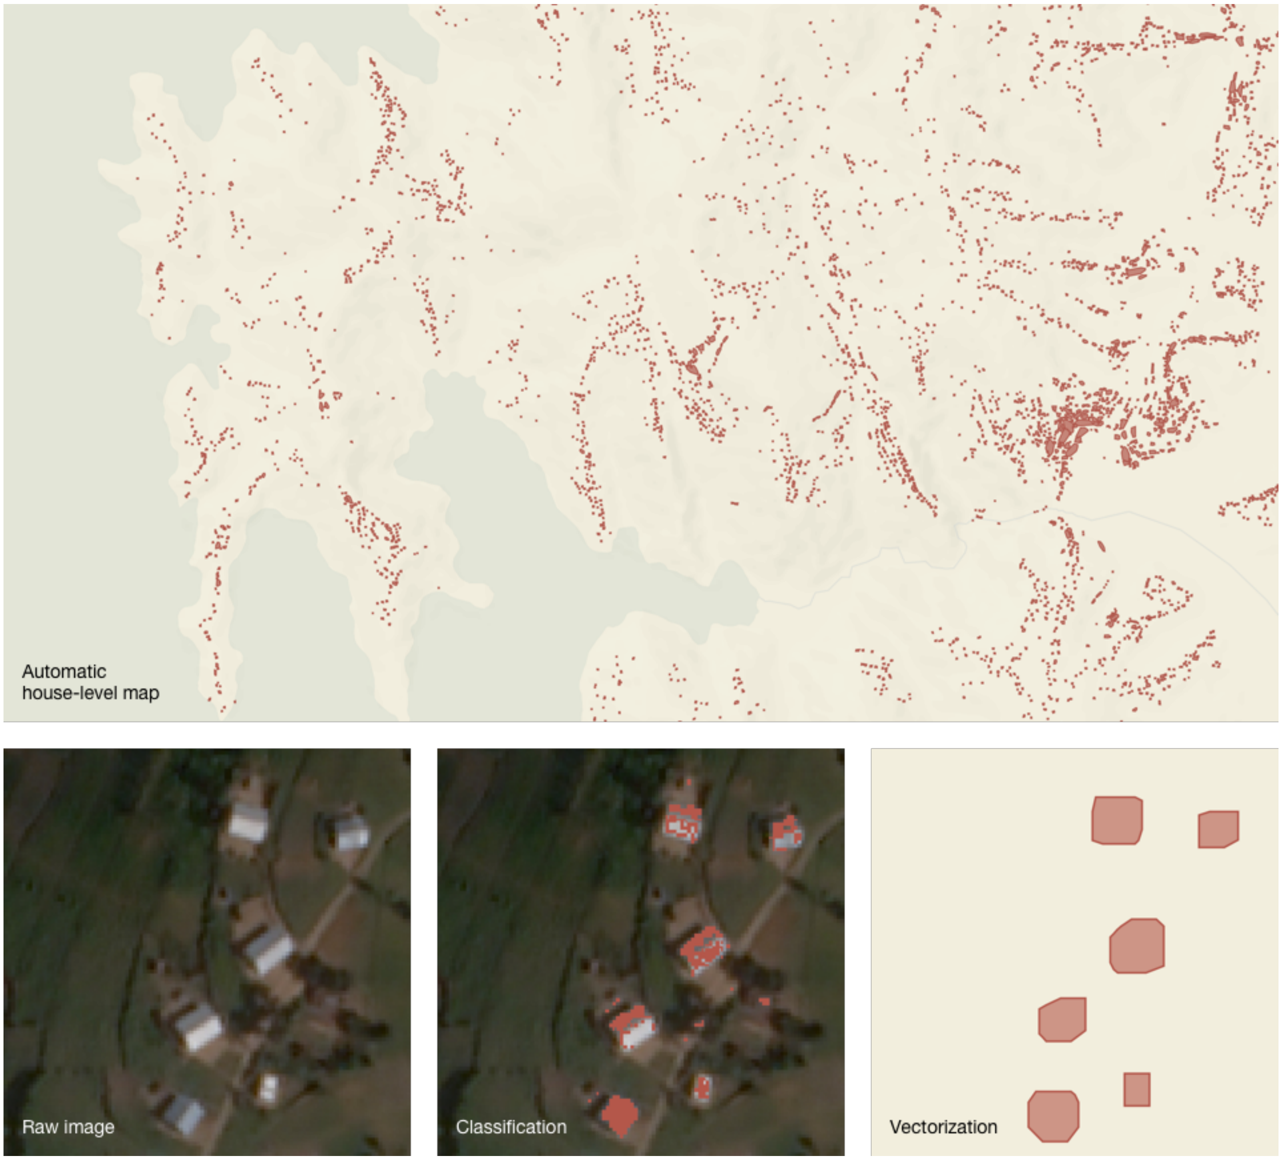

Figure 3.5: Automatic generated house-level map from raw satellite imagery using machine learning and image processing.

(Breiman 2001) as implemented in Earth Engine:

Step 1 Add the pixels in an area of 2.4 m (3 × ground sample distance) around the annotation markers to the training data.

Step 2 Create the training dataset by merging the manual annotations of "building" and "not building."

Step 3 Train the two-class random forest classifier with the labeled dataset and 100 decision trees on all five spectral bands.

Step 4 Classify all pixels in the image as "building" or "not building" using the trained classifier (cf. fig. 3.5, bottom middle).

Step 5 Remove noise in the classification through morphological opening by (1) erosion of 1.6 m, (2) dilation of 2.4 m.

Step 6 Convert homogeneous building pixel regions to vectors at the native pixel resolution of 0.8 m.

Step 7 Simplify each polygon by computing a convex hull around it (cf. fig. 3.5, bottom right).

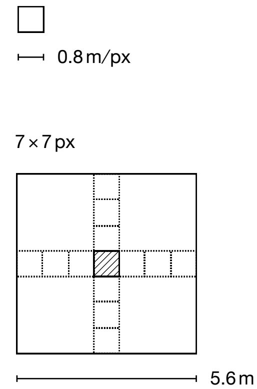

Random forest classifiers require few hyperparameters other than the number of trees to use. Generally, the more trees, the better the result, but also the more computationally expensive to build the forest, especially regarding of used memory. Because our analysis runs on cloud computing infrastructure, this is less of a concern. In most cases, the predictions become better with more training data. By using single-point annotations as input data (in contrast to relying on building footprint polygons to train the model) we are well-positioned to potentially merge ground annotations from GPS data into the training set.

3.4 *Mobile App Design and Engineering*

User Interface Design

Our approach in designing the mobile app and activity prioritized the community and built upon the insights from our field observations (section 3.1). We designed the system involving the community's perspective through human-centered design. This creative approach to interactive systems development has been successfully employed for projects in the social sector (IDEO 2015). We specifically created a design that empowers CHWs to understand the collected data and share it with their community.

We sketched solutions using pen and paper (fig. 3.6), designed mockups in the UI design tool Sketch (fig. 3.7), created prototypes for user testing in Keynote, and built functional prototypes of app elements natively on iOS. The design process was iterative, often turning from writing code back to sketching.

The mobile app is named *Yego* (from Kinyarwanda, "yes"). Its interface centers around the CHW, their location, and their household catchment. The UI, exhibited in fig. 3.6, has two main modes to view this data, the list view, and the map view. The app launches into the list view, and all tasks can be accomplished from this screen.[25] A user profile view at the top shows the user's photo and name, their village and catchment size, and a settings button. The settings

[25] The counterintuitive insight of centering a mapping app's UI on a list instead of a map came from observations during map sketching with CHWs. The resulting maps suggested that understanding cartographic representations is an acquired skill and hence would not serve well as the primary way of organizing their geographic catchment information.

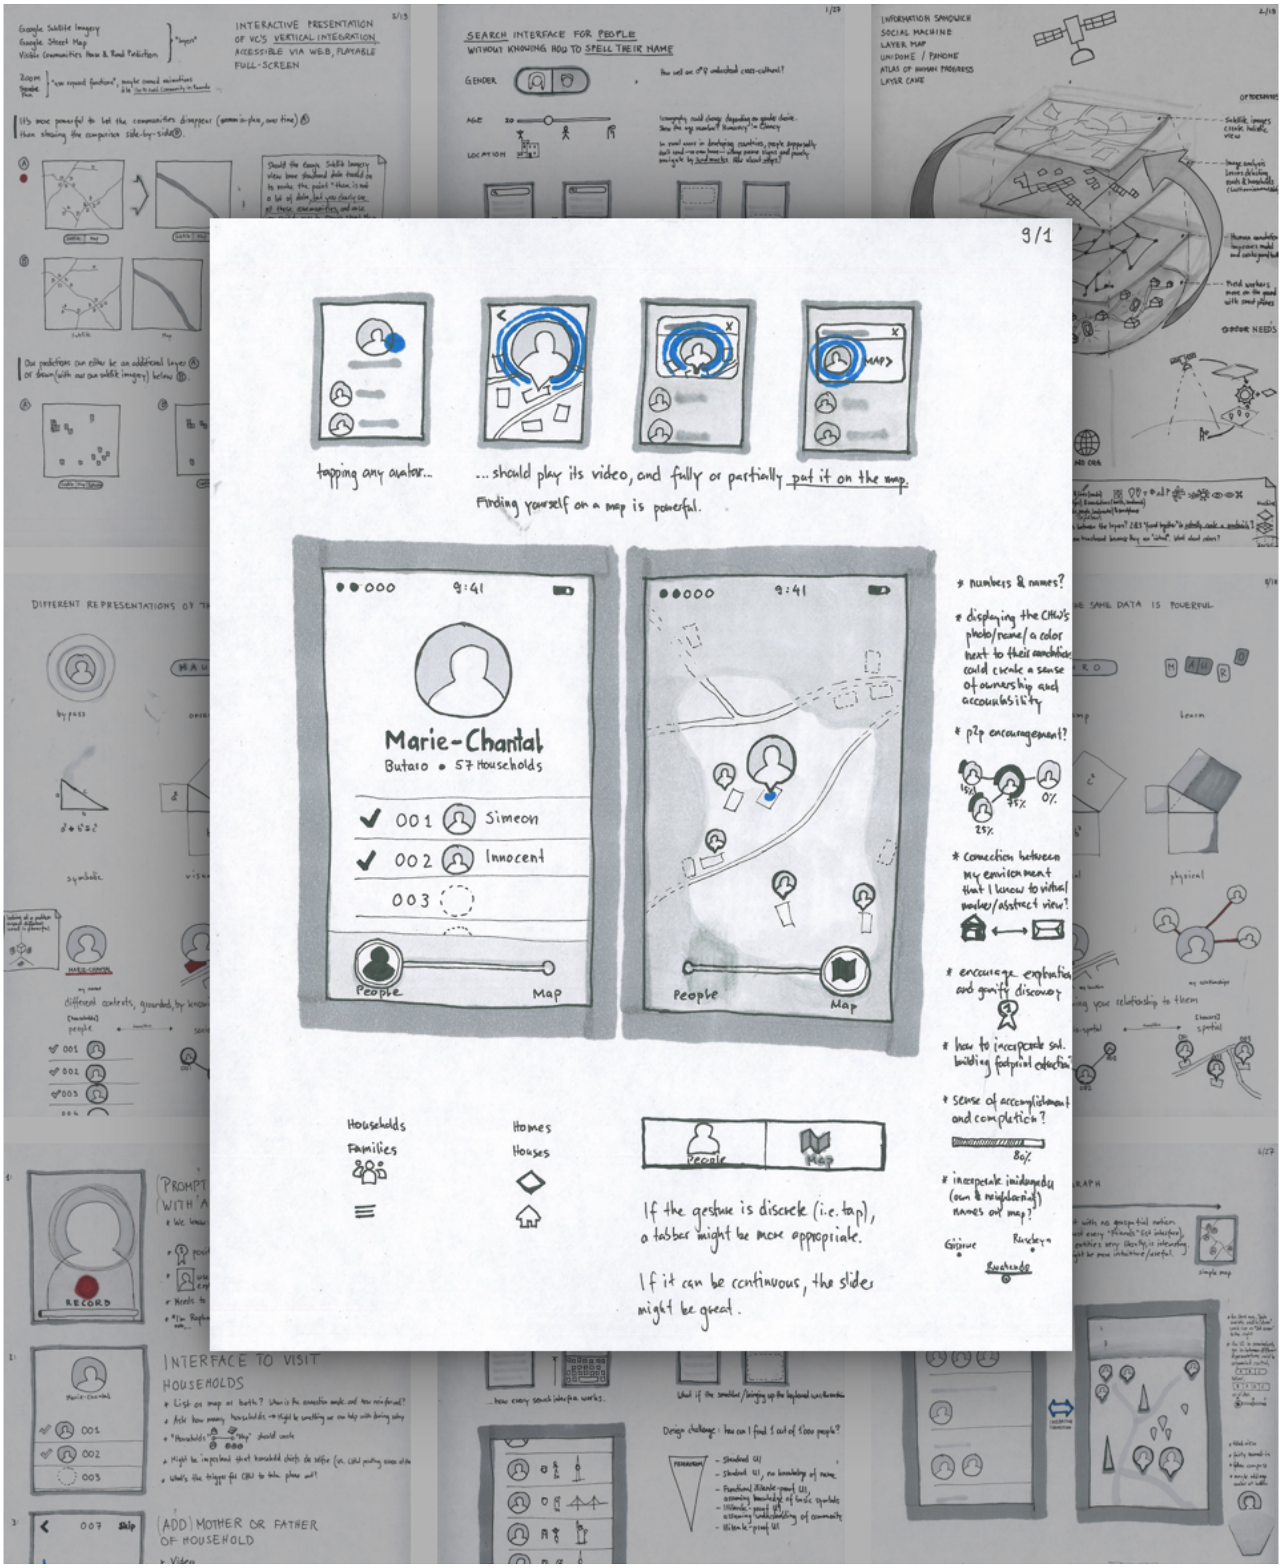

Figure 3.6: Sketches during the design process. Human and machine aspects, as well as social and spatial aspects, influenced an initial wireframe for the mobile app interface.

screen enables configuration of the app during setup, and change settings during use. The list of households is prepopulated with a numbered row per household. There are as many rows as there are households in that specific catchment. Mapped rows show a filled circle around the number, a timestamp of the check-in, and a checkmark. Tapping the checkmark enables the user to uncheck a check-in. A confirmation dialog prevents removing a check-in by accident.

There are three ways to transition to the map view: (1) mapping a new household by tapping on an unmapped row, (2) locating a mapped household or oneself by tapping any circular shape, (3) swiping the list view to the left of the screen to reveal the map behind it.

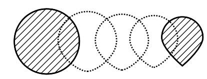

The list view remains partly visible on the left screen edge, and users can still fully interact with it. Interaction with the list from the map view enables the mapping of households and locating of already mapped households even from this view. Swiping to the right on the list, transitions back to the full list view.

The app uses intuitive design elements, references elements that the user is familiar with from their environment, and uses as little text as possible. Certain elements that are best expressed textually or numerically are represented with short labels. Such labels are translated into the local language and format (e. g. date format).

Satellite imagery is another central design element in the app. The map view can be navigated by moving (swipe or pan) and zooming (pinch in or out). Although this region is largely unmapped, we leverage the imagery in the map view to show the user a visually rich and relatable basemap. Showing a cartographic map of the area would lead to a mostly blank screen due to the lack of available mapping data. The machine-generated building footprints are overlaid on the map. Mapped locations are displayed as markers. Where markers intersect with a building footprint, that area's color turns from red to green.

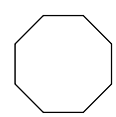

For our design to feel empowering and motivating to the CHWs, we employ gamification techniques. Such elements are not added-on but part of the core experience. The gamification framework Octalysis[26] describes elements present in our design:

[26] Chou (2016) introduces an octagonal-shaped diagram arranging eight core human drives. Drives on the left are extrinsic, motivating us towards a goal with reward. Drives on the right are intrinsic, we do it out of joy, and no reward is necessary. The bottom has negative drives that are not sustainable long-term, and the top positive ones, like greater meaning.

OWNERSHIP of the user over the app is given by showing the user's photo, acting as an avatar. This avatar is displayed in the profile area at the top of the screen and again as an indicator of the current location on the map.

ACCOMPLISHMENT of completing a task, like quests on a list, is ex-

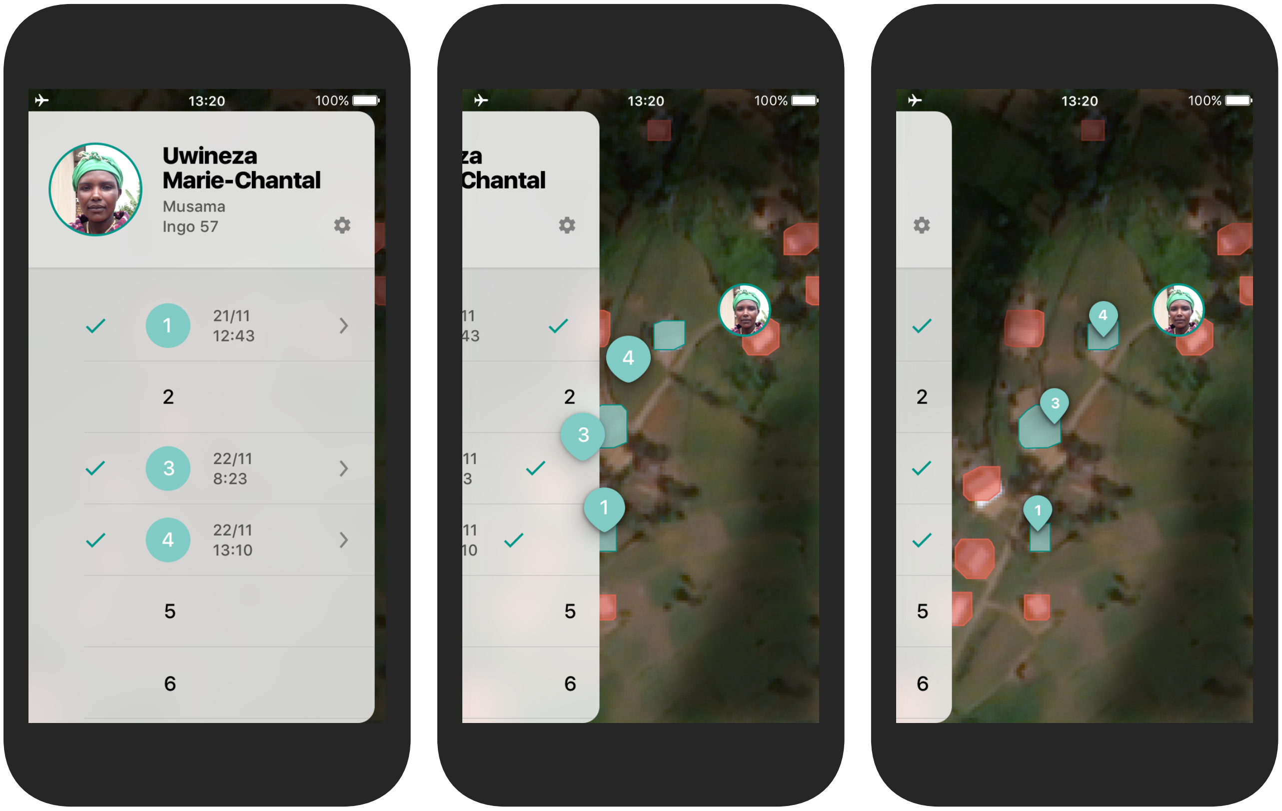

Figure 3.7: The design of our mobile app Yego resembles CHWs' familiar paper-based list (left). Transitioning from list view to map view (middle) locates the households geospatially on a satellite map that includes a layer of automatic building footprints (right).

perienced when checking off rows on the household list. The checkmarks demonstrate achievement symbolically. The accumulation of complete rows acts as a progress bar.

Unpredictability is reinforced by lifting a fog overlay and revealing the map where the user physically explores their environment. The exploration is directed and further encouraged by the highlighted building predictions shining through the fog like collectibles.

Empowerment to help shape the map is provided to individuals when their actions directly result in changes to the map. When a household is mapped, instant visual feedback is given. After the loading animation completes, the view changes to the map and a marker drops to the current location, turning the color of detected buildings from red to green.

Meaning and a sense that the user is part of something bigger is provided by displaying the size of the user's catchment area and the name of the greater village. These elements appear in the profile and build a personal and communal narrative.

In summary, the process that led to this design is based on insights from the community from our field observations discussed in section 3.1. We then further iterated on the design back at the MIT Media Lab. The design characteristics we derive from evaluating its performance during the field pilot (section 3.6) are discussed in chapter 5.

Software Architecture

The app was built natively for iPhones running the latest iOS 10. We chose iOS for most effective rapid prototyping, drawing from our experience and leveraging the platform's mature software development kit.[27] With some additional effort, a comparable architecture could be replicated for other smartphones, such as Android devices.

[27] In addition to the frameworks contained in the iOS software development kit the app leverages mostly geo-oriented functionalities from the open-source frameworks `iOS-GPX-Framework`, `iOS-KML-Framework`, `GeoJSONSerialization`, and `MBProgressHUD`.

We implemented the app in Objective-C using Apple's Xcode development environment.[28] The app's main software design pattern is Model-View-Controller. This pattern assigns roles to each object and how they communicate with each other. Models encapsulate the data and some basic behaviors, views present the UI, and controllers tie the model to the view (Apple 2015). We employ this pattern in every screen for reusability and extensibility of the objects. The main entry point in the app is a root controller responsible for bringing the list, map, and settings view controllers on the screen and mediating data flow between them. The map view controller ties an additional seven map-related classes together that extend the functionality of system frameworks. The model layer consists of a class for the user (singleton) and one for households.

[28] The codebase consists of 3591 lines of code (2252 excluding comments and blank lines) across 17 classes.

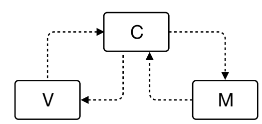

The three core aspects of the software architecture – the high-performance offline satellite image basemap, acquiring location using GPS, and client-side offline data persistence – are described in more detail in this chapter.

We implemented all planned features, and the app is fully functional. The UI fully matches the design specification, with two differences to stay within the budgeted implementation time and better app performance: (1) transitioning between list and map view is achieved with a bottom tap bar instead of swiping, (2) no fog is covering the satellite map.

Satellite Image Basemap

In addition to using the custom acquired imagery for building prediction, we describe in the interface design section why a satellite image basemap is a central visual design element in the app. Because our app must function offline, this data has to be bundled with the app. This is challenging with the limited processing, memory, and storage capacities of smartphones compared to the cloud computing infrastructure we used to work with the data thus far.

In section 3.2 we establish that high-resolution imagery, even of a small geographic region, results in huge files. To browse maps of the entire world at high resolution, mapping services rely on tile pyramids[29] to load only the viewed part of the image and only at the minimum resolution needed to perceive a crisp image at the current zoom level.

[29] Analogously to how the image was stitched together from multiple smaller images when it was acquired by the satellite, tile pyramids divide the image into a grid of smaller manageable tiles. Repeated for different zoom levels at different resolutions, this results in a pyramid-shaped structure. In-between the levels, tiles are interpolated for smooth zooming.

At zoom level 0, the entire world displays in a single 256×256 pixel tile. At level 1, while keeping the tile size a constant 256×256 pixels, we double the tile horizontally and vertically to a total of four tiles and a map of 512×512 pixels. At level 2 the map consists of 16 tiles and is 1024×1024 pixels big. This is repeated for each level as shown in table 3.2.

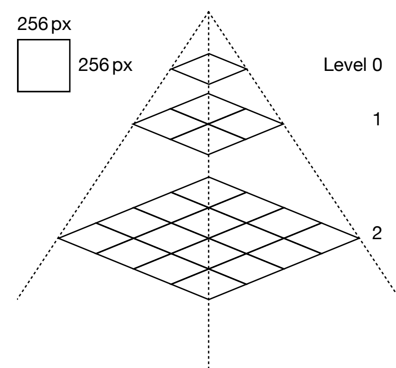

The number of required tiles per level is calculated with $2^{zoom} \times 2$. Base 2 raised to the power of *zoom* to double at each level, multiplied by 2 because the map is two-dimensional. The maximum *zoom* is 20 for most practical applications.

The spatial resolution S in meters per pixel is calculated for each level with $S_{zoom} = C \times \cos(\phi) \div 2^{zoom} \div 256\,\text{px}$, where C is the (equatorial) circumference of the earth (40 075 016.686 m) and ϕ is the latitude in degrees (0 at the equator).

As table 3.2 shows, the tile pyramid grows exponentially and storing a single zoom level with standard JPEG image compression requires giga- or even terabytes. Our image has a resolution of $\sim 0.8\,\text{m/px}$ so our maximum level is 18. To estimate the storage requirements of our offline tile pyramid, we assume a pilot site of $10\,\text{km}^2$:

$$\sum_{zoom=1}^{18} \left(\frac{10\,000\,\text{m}}{S_{zoom} \times 256\,\text{px}} \right)^2 \times 40\,\text{kB} \approx 235\,\text{MB}$$

We divide the region size by the size one tile covers, square the resulting number of tiles for the two dimensions, and sum the tiles up for all 18 zoom levels. The resulting hundreds of megabytes can be bundled on a smartphone.

We used Earth Engine with its powerful computing infrastructure to pre-generate the tile pyramid. To render the map view, the app uses Apple's MapKit framework.[30] The translation between geographical coordinates and zoom level to tile indices is calculated by MapKit. Each tile is represented by an object loading the data from the device's non-volatile flash memory to render the raster image. The building footprint vector data and household location markers are added as an overlay on top of the basemap and rendered on demand for any given map extent to maintain a responsive experience.

[30] MapKit was introduced with iOS 3 and is typically used by apps to display standard maps loaded on demand from the Internet. Among other functionality, the framework handles basic gestures such as pan and zoom and can display a user's location.

Table 3.2: Ground resolution, number of tiles, and estimated storage requirement for 20 zoom levels of a typical mapping tile pyramid.

Zoom	Resolution* (m/px)	Tiles	Storage†
0	156 543	1	40 kB
1	78 272	4	160 kB
2	39 136	16	640 kB
3	19 568	64	3 MB
4	9784	256	10 MB
5	4892	1024	41 MB
6	2446	4096	164 MB
7	1223	16 384	655 MB
8	611	65 536	3 GB
9	306	262 144	10 GB
10	153	1 048 576	42 GB
11	76	4 194 304	168 GB
12	38	16 777 216	671 GB
13	19	67 108 864	3 TB
14	9.6	268 435 456	11 TB
15	4.8	1 073 741 824	43 TB
16	2.4	4 294 967 296	172 TB
17	1.2	17 179 869 184	687 TB
18	0.6	68 719 476 736	2749 TB
19	0.3	274 877 906 944	10 995 TB
20	0.15	1 099 511 627 776	43 980 TB

* Values for projection at the equator
† Assuming an empirical average tile file size of 40 kB

Location Acquisition

To determine the current location, smartphones commonly combine different sources such as GPS, cell tower trilateration, Wi-Fi hotspot databases, and even motion sensors and vector mapping data.[31] While GPS generally is the most accurate, it has the highest acquisition time and power consumption. The only source that works offline though is GPS. Fortunately, GPS signals are good in rural areas because there is no occlusion by urban street canyons. Furthermore, the conditions to acquire a GPS location close to the equator are the most optimal. With the GPS chips in modern smartphones, the accuracy is to the meter.

[31] A main source for assisted GPS (A-GPS) location acquisition is the location of nearby cell towers and Wi-Fi hotspots. Through signal strength to available stations the device can measure the distances and with trilateration (triangulation, in contrast uses angles) determine the location.

The app uses Apple's CoreLocation framework to acquire geolocations.[32] We configured the framework for highest accuracy within the given constraints of no connectivity and no existing validated mapping data to incorporate in the calculations. When mapping a household, the app requests a single point location[33] in a background process. Upon acquisition, the app determines its horizontal accuracy, which can vary depending on the physical location of the device

[32] CoreLocation was introduced with iOS 2 in 2008 but the first device with a GPS chip was not released until 2010 with the iPhone 3G.

[33] It would be valuable to have GPS coordinates of each corner of the building, but that technical requirement would require a significant change in the human activity.

(e. g. inside a building compared to under clear sky). If the accuracy passes a certain threshold,[34] the app stores the location. Otherwise, it visually notifies the user and enables them to retry.

[34] We determined everything above ± 50 m as insufficiently accurate. Our tests measured most GPS fixes with the technically highest supported horizontal accuracy of ± 5 m.

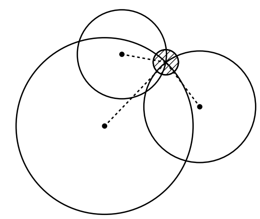

We considered power consumption for all our engineering decisions and determined continuous background tracking of the location as too power intensive. However, whenever the app is in the foreground, and the GPS is enabled in any case, the app saves all locations that pass the accuracy threshold. Each location includes latitude, longitude, horizontal accuracy, horizontal speed, altitude, and timestamp. The app combines those location points in a track and saves them to the device's non-volatile flash memory in the standard GPX format. A setting allows enabling and disabling of this feature during runtime.

Data Persistence

The app is used in a disconnected environment, and thus the accumulated data must be stored locally on the devices without cloud backup. With that consideration, we engineered the data persistence layer to minimize the risk of data loss to the physical loss of the device.

The first measure for flawless local data persistence is a robust way to serialize the data and to write it to the device's non-volatile flash memory. Each model object has straightforward instructions how to serialize and deserialize itself. Using these instructions, object graphs are written out to the flash memory as a binary archive. This happens on all critical events during the app's lifecycle, such as mapping a new household and putting the app in the background. Upon app launch, the app loads these archives to restore its state. It keeps user profile data, user location information data, and household mapping data in separate archives to spread points of failure.

Another measure is diligent error handling in case of system failure or data corruption. Each potential source of error is handled individually by trading off chance of data recovery with risk of continuous errors. For example, while it is standard practice for connected apps to simply reload data from the server in case of a client error, our app attempts to protect the existing data by automatically terminating itself and enabling the user to retry and load the data into memory upon a subsequent app launch.

Furthermore, the app has local persistence redundancy. In addition to the core persistence layer, the app writes the data to different standard formats (JPEG, JSON, KML, GPX) that are appropriate to their content. This also allows for in-field exporting and testing of the data.

3.5 Field Pilot Hardware

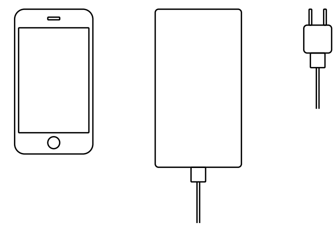

In addition to fulfilling the technical requirements, we set out that the selected field hardware had to be appropriate for a low-resource setting without compromising on user experience. The smartphone had to be protected for field use, and a non-intrusive way to keep the phone charged had to be found. We sourced twelve field kits for the pilot participants, plus three extras as backup. Each CHW received a kit consisting of:

Smartphone We selected the low-priced iPhone 5C to run the latest version of iOS. We purchased refurbished models to target the under-$100 category of comparable low-end Android devices. It uses a low-power GPS chip.[35]

[35] Part of the radio frequency transceiver Qualcomm `WTR1605L`.

Portable battery pack In addition to the primary intended grid charging, we provided an AmazonBasic Portable Power Bank. We chose the capacity[36] to allow for five days of off-grid use. We evaluated[37] and discarded solar options[38] for our purposes.

[36] $capacity_{pack}\ 10\,\mathrm{A\,h} \div (capacity_{phone}\ 1.510\,\mathrm{A\,h} + 20\,\%\ loss) \approx 5$

[37] $(capacity_{phone}\ 1.510\,\mathrm{A\,h} + 20\,\%\ loss \div 6\,\mathrm{h}\ sunshine) \times 5\,\mathrm{V} \approx 1.5\,\mathrm{W}$. This means for daily solar smartphone charging a sizable panel that would be best installed on the roof is required.

[38] We evaluated solar lanterns using MIT D-Lab's Off-Grid Energy Group's product comparison. However, in addition to introducing higher cost and complexity, we were cautioned about the implications of selectively introducing lighting in a community.

Accessories A protective case and a locally compatible wall charger.

We configured[39] each device to only require the brief in-app user setup during deployment in the field.

[39] We installed our app Yego directly from the development environment and provisioned it with our developer account, put restrictions in place to protect from accidental deleting of the app and its data, and set the locale to Rwanda and the preferred language to Kinyarwanda.

3.6 Field Pilot in Rwanda

Preparations and Plan

We determined that running a small pilot with representative target users in their environment is the most effective method to test the feasibility for a community to self-map using our system and observe its effect. From what we have learned during our field observation (section 3.1), working with PIH-supported CHWs in one of the rural districts is feasible and the gathered evidence significant. Strong results from this group would suggest similar communities at the BOP could use it too. We were invited by PIH to run a field pilot at their Burera site, and they facilitated obtaining the required local government authorization. Our study design including interviews was reviewed and approved by MIT's Internal Review Board.[40]

As the pilot site, we chose four representative villages in the Kaganda cell, Burera district. They are rural communities with PIH-supported CHWs and are accessible within a day from the Butaro Hospital.[41] All CHWs in that region were offered to participate in the study voluntarily. PIH facilitated contacting the twelve potential participants and delivering a translated version of the consent form

[40] MIT investigators conducting research involving human subjects complete additional training and have the Committee on the Use of Humans as Experimental Subjects review and approve all materials of the study.

[41] Our field base was the PIH staff operations at the Butaro Hospital. For access to off-road-capable cars with drivers, accommodation with electricity, and meals, there was no alternative infrastructure available in that area.

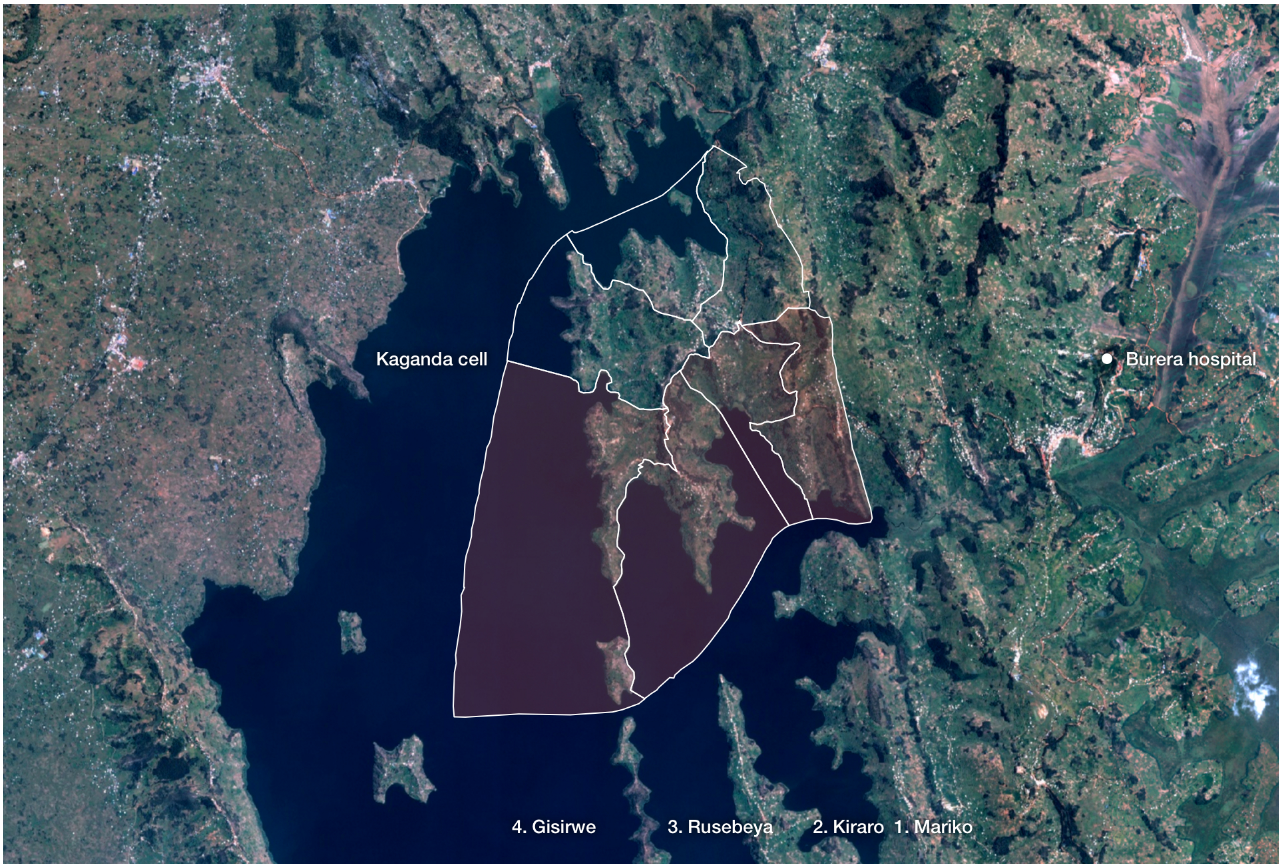

Figure 3.8: Map of the pilot site with its four selected villages in the Kaganda cell highlighted (in order of approach from the hospital): Mariko, Kiraro, Rusebeya, Gisirwe

and description of the activity. Their role as CHW requires them to be literate in the local language. The form describes the activity, reimbursement for expenses,[42] data privacy, and that the device should be returned at the end of the study. All twelve CHWs were interested and willing to participate in the study.

[42] While participation in the study was unpaid, the form states two reimbursements: (1) additional time participants voluntarily took out of their day for interviews (at an appropriate rate as advised by PIH), (2) the paid charging of the smartphone at the nearby commercial center (at the local market rate).

Participating CHWs gave informed consent for the collected data to be used publicly so the map, for example, can be published and distributed. In the consent form, participants could additionally choose to be attributed by name and photo, which all participants chose. The identity of served households is coded, and the CHWs only record locations. Patient and clinical data remain in CHWs personal hardcopy records.

The pilot duration was set for ten days to allow for a natural cycle of routine household check-ins by the CHWs. We setup all required hard- and software at the MIT Media Lab and brought it to the field.[43] The plan was to start the study with six CHWs on day one by setting them up with the tool at their homes and do the same with the remaining six the following day, offsetting the schedule of

[43] To ship equipment for field work, our mentors' advice, and own experience is to bring it in the cabin luggage when traveling. It is the safest and easiest way in most cases – even for unusual devices or quantities. To minimize risk, we carried an official letter with a complete list of equipment and its intended use.

that cohort by one day. This was determined by the constraints of our travel time to the communities and estimated duration for the initial individual setup. Then we would leave the CHWs with the devices and not return until the respective fifth day. On those days, we would help with any potential technical problems and document the activity by taking photos. This visual ethnography is intended to capture context and to aid in understanding users' environments. On each CHW's tenth day we planned to collect the data and hardware. In the end, the devices would be left with the PIH research department.

Field Work

Before starting any activities in the field, we met the local authorities and used the opportunity of a monthly gathering called *umuganda*[44] to introduce ourselves to the entire community beyond the recruited participants. The logistics for the over 30 field trips for setup, check-in, and data collection required significant coordination between CHWs' availabilities and the required PIH resources.[45] We were accompanied by PIH staff for all field visits and employed a translator for all interactions with the CHWs. The interpreter helped translate the language and more importantly, the nuances of the cultural context.

During our first visit to the CHWs, we conducted an initial interview, asked them to sketch a paper map of their mental image of the village using the method Lynch (1960) describes and provided them with the field kit. In individual training that lasted approximately 15 minutes the CHWs learned the basic use of the hard- and software. Topics included how to handle the device and keep it charged, how to interact with the touch screen and basic usage of the phone, and how to integrate the mapping into their routine activity.[46] Together, we configured their profile in the app during this setup. During our final visit, we conducted an exit interview, repeated the sketch map exercise, and collected the data and hardware.

The direct human engagement was critical to test our bottom-up community approach. We helped to integrate the solution into participants' regular workflow and guided them to explore other functions of the device. We also incorporated the principle of giving the data back to the community into our field work by bringing a portable printer and creating maps from the data while still in the field. The process and results of leaving an understandable representation of the participant's activity are described in section 5.3.

[44] The national community work day *umuganda* takes place every last Saturday of the month. Citizens improve eroded roads or build homes for vulnerable people. It is mandatory to participate, and people embrace this duty. That Saturday we carried stones from a quarry to the village center for a new school.

[45] For all field activities, it was critical to account for buffer times and reducing critical path dependencies.

[46] We used role play at CHWs' homes to simulate a home visit. The steps were: (1) walk to a home in your catchment, (2) launch Yego and look for the number assigned to that household, (3) position yourself at the entrance and check that number off. Because the app reacts to the user's location, physically demonstrating it was the most effective training.

Figure 3.9: Impressions from the field pilot: setup at a CHW's home (top left), our local collaborator and translator introducing us to the local community (bottom left), and CHWs mapping in the field during routine check-ins (right).

Data Collection

Because the app functions entirely offline, the accumulated geospatial data was stored solely on the devices during the pilot. Upon completion of the pilot, we manually collected the data by moving it off the devices and backing it up. We did not collect any sensitive data that would need to be secured. During our field visits, we acquired GPS locations of the twelve CHWs' homes ourselves as reference points for data validation.

We conducted all interviews with the same questionnaire and took structured notes for all observations. The initial interview captured background information and the exit interview captured information about the app usage and what might have made the activity easy or hard.[47] This information was collected to use as evidence when assessing UI decisions and the effect of the tool's usage on the users.

[47] The questionnaire included open-ended qualitative questions and quantitative questions on Likert scales. The questions can be found in appendix A.

4 Results and Analysis

In this chapter, we present and analyze the data that twelve CHWs collected during the pilot of our self-mapping system. We use mixed methods and additional survey data to evaluate the system's performance, the effect of our mobile mapping app Yego on participants, and the derived socio-spatial maps. The results strongly indicate that with an appropriately designed solution, communities at the BOP can create state of the art maps. Without prior knowledge of smartphones, the participants mapped 100 percent of the homes at high accuracy in the study's region. Participants found the app Yego easy to use due to the machine assistance and the alignment with routine community health tasks. The design of Yego enabled participants to develop new spatial insights of their communities that are useful for their job. Our spatial data analysis yields high-resolution population density maps and spatial CHW catchment areas that did not exist before for this region. We further combine the spatial with social data and model a socio-spatial community network and map. Our analysis yields measures for social centrality and the map representation enables observations of the correlation between a household's social centrality and its spatial location in the village.

4.1 Demographics of Participants

The participants of the study were twelve CHWs from four villages in Burera. We interviewed the CHWs during the pilot and obtained demographic information (table 4.1). For easier reference, we assigned a number to each village (1–4) and numbered the three CHWs within each village. For example, for village "1" CHWs were numbered "1.1", "1.2," and "1.3."

There is an even ratio of males to females among our twelve CHWs. All are deeply rooted in their communities with an average of 25 years lived in their villages and four born in the same village they live in today. On average, they have served as CHW for almost ten

	μ	σ	1.1	1.2	1.3	2.1	2.2	2.3	3.1	3.2	3.3	4.1	4.2	4.3
Gender	—	—	♀	♀	♂	♂	♂	♀	♀	♂	♂	♀	♀	♂
Age	40.1	8.2	33	45	34	40	49	33	55	53	41	30	33	35
Size of household	6.3	1.6	6	7	6	8	6	6	5	10	8	5	4	5
Years in village	24.8	15.7	11	20	13	40	49	10	34	53	14	10	9	35
Years of schooling	6.3	1.2	6	8*	6	7*	9*	5	6	4	6	6	6	6
Years as CHW	9.1	5.1	10	9	10	5	20	8	10	17	3	2	5	10
Size of catchment	52.7	13.5	82	70	68	46	45	44	40	38	38	55	51	55

Participant identifier format: `village.chw`
μ = mean, σ = standard deviation
♀ = female, ♂ = male
* Includes secondary schooling

Table 4.1: Demographics of the twelve CHWs ($N = 12$) that participated in the study.

years. The average catchment size is 50 households with a broad range from 38 to 82. All CHWs have attended at least four out of six primary school years and hence can read and write in the local language Kinyarwanda. The three that have attended a few years of secondary school were the ones that demonstrated a little spoken French.

All CHWs are married with children and work as subsistence farmers living among the rest of the community. None of them have electricity or running water in their homes. While all CHWs have a mobile phone, none of them had previous exposure to a smartphone.

4.2 *Collected Dataset and System Performance*

The total number of households across the twelve catchment areas is 632, including the CHWs' own households. Within the ten days of the study, each CHW mapped 100 percent of the homes in their catchment area, resulting in complete map coverage of the study's region. With an average household size of 4.4 people in rural Rwanda (National Institute of Statistics of Rwanda 2012, p. 25), the estimated population of our pilot site is 2781. The map in fig. 4.1 shows the 632 buildings mapped by the CHWs. The only available mapping data[1] before the pilot was the major road in the northeast portion of the map (top right in fig. 4.1).

[1] The mapping data is by OpenStreetMap. While no popular mapping service has a single building mapped in this region, the level and quality of mapped roads differs among services.

We collected the data directly from the smartphones that the CHWs used because the app operated completely offline. CHWs participating in the study returned all twelve smartphones in undamaged conditions. The app ran crash-free, and the locally stored databases (as described in section 3.4) remained robust, resulting in zero data loss.

We evaluated the acquired coordinates qualitatively while still in the field. Using the app, we zoomed into the high-resolution satellite

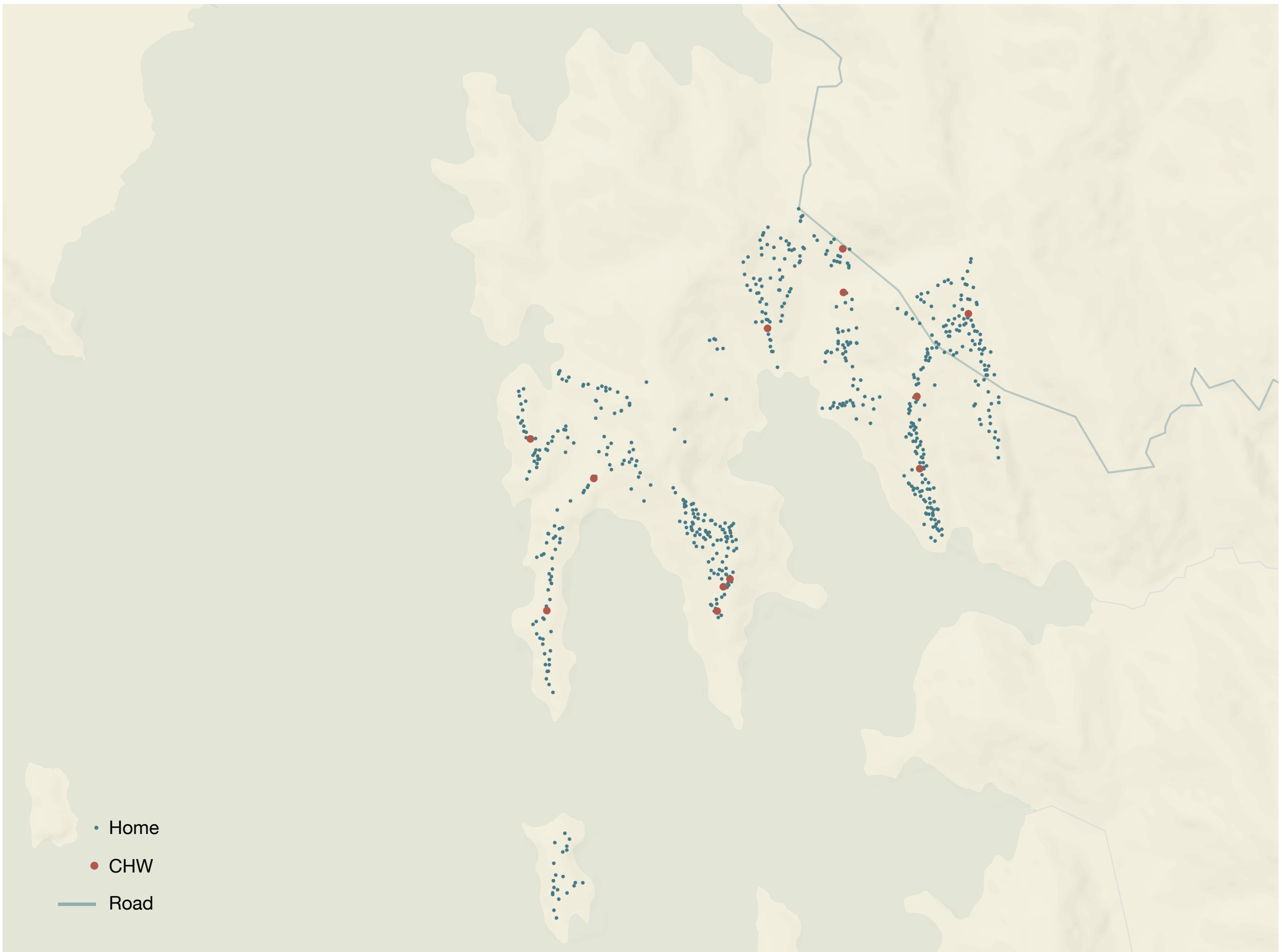

Figure 4.1: Map visualization of all 632 homes mapped by twelve CHWs.

imagery in the map view and visually assessed the mapped locations relative to building footprints (a process we observed the CHWs doing themselves as well, which might have led to corrected mistakes and thus increased data quality). The accuracy was strikingly high. We confirmed this result after aggregating the data and reviewing it on a laptop computer.

The quantitative data on the accuracy of the acquired coordinates reflects the qualitative assessment: the horizontal accuracy ranged from ±5 m to ±50 m with a mean of ±8 m. The technically highest supported accuracy of hard- and software is ±5 m. Over 97 percent of all coordinates were acquired in the ±5 m to ±10 m range. Each coordinate consists of a longitude, latitude, and altitude.[2]

[2] Example coordinate:
```
[29.78607469241115,
  -1.417631702498922,
  1968.255859375]
```

The collected dataset for our analysis contains social and geospatial data. Table 4.2 shows which actor knows or has access to what elements of relevant data. While the CHWs know all households

	Household name	Household identifier	House coordinates*
CHW	knows	collects	collects
PIH	knows	knows	accesses
MIT	—	accesses	accesses

* Novel data

Table 4.2: Elements of household- and house-level dataset (social and geospatial) with access rules to protect privacy.

in their catchment by name, those catchment lists are not digitally stored. By capturing the catchment sizes with the app, we can derive identifiers for households. For example, household number "7" of CHW "1.1" is assigned the identifier "1.1.7."[3] We did not collect any personally identifiable information of the community *members* to protect their privacy. The home locations were previously only known to CHWs and had never been recorded. During our pilot, the CHWs added this novel element to the dataset.

[3] Household identifier format: `village.chw.household`

All participating CHWs received brief training to use and maintain a smartphone, and use the installed app Yego. During the study, CHWs could reach us using their regular mobile phones. We received calls from five CHWs during the first two days reporting their progress and asking minor questions (e. g. increasing the catchment size in the app to account for a new family). It is worth noting that all calls were from CHWs we set up on the first day of the study. A possible interpretation of this difference between the first cohort and the second cohort that was set up one day later is that this is a sign of peer-to-peer learning. That is to say, already experienced CHWs of the first group may have helped CHWs of the second group without intervention from our side. Using the CHWs' answers to questions from our survey,[4] we evaluate how our system performed with respect to the CHWs' use.

[4] The survey questions are in appendix A and section 4.3 describes how the interviews were conducted.

Although none of the participants had previously used a smartphone, they all found the app Yego easy to use. On a difficulty scale from 1 to 5 in which 1 is "very difficult" and 5 is "very easy" they rated it at 4.2. They mentioned the few required steps as the main reason for the ease of use. One participant said, "It's just one touch from me, and everything else happens automatically." (CHW 2.2, *Field Interview* 2017) An important factor for the activity to be well integrated in CHWs' days was that they could combine it with regular community health activities and did not have to wait for a head of household to be home to complete the task. They reported an average of 40 percent of head of households home when they visited their homes.

First-time handling of a smartphone is generally challenging, and the CHWs acknowledged the initial lack of experience. For example, one CHW had difficulties unlocking the phone, and another one found scrolling to be a challenge. However, they quickly adapted to the new experience and found our minimal training as described in section 3.6 at the beginning of the study helpful. Once they were familiar with the smartphone, they all preferred it over their regular mobile phones. All participants mentioned that functions were easier to see and quicker to access on the touchscreen compared to navigating the text menus on a small display in a foreign language.

The CHWs liked using Yego. They appreciated that the list had all their households; the clear visual indication of which ones were mapped, giving a sense of completion; being able to correct mistakes; and that their data was still there when returned to the app. Because CHWs are familiar with the household numbers and to which home each number belongs to, they mentioned being confident in using the list and did not have to search for the location. They also liked the smooth transition from the familiar list to the increasingly familiar map and were delighted to find the same numbers there.

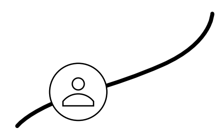

Although initially unfamiliar with the map, we collected many positive comments about interactivity with the map. Although the map does not serve a direct functional purpose in the app, we describe in section 3.4 its choice as a central design element for Yego. CHWs liked how their avatar photo followed them in real time on the map and were surprised that the paths they were walking on were visible without having to add them to the map. Most CHWs reported a strongly positive experience when they realized that this view corresponds to reality. One CHW noted, "It's the real world, and I can see myself and the buildings as I know them in it." (CHW 4.1, *Field Interview* 2017)

Participants used the smartphones beyond the mapping activity on average eleven times per day. All CHWs used the camera to take photos and videos. Three reported continued use of Yego after mapping to explore the map and look up mapped households.

CHWs adopted the usage of Yego without significant complications during the duration of the pilot. The resulting user engagement, data completeness, and accuracy exceeded our expectations. The collected house- and household-level map establishes the foundation for the social and spatial maps as introduced in chapter 1. The mapped households are the social nodes and the locations of their homes the spatial nodes as shown in fig. 4.2. Later in this chapter, we further analyze this data through different lenses and explore how links between these nodes can be established.

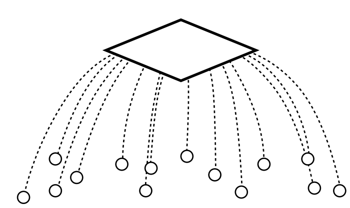

In summarizing, the results strongly indicate that with an appropriately designed solution, communities at the BOP can create state of the art maps. In the next section, we analyze the direct impact on CHWs from engaging with the app during this study.

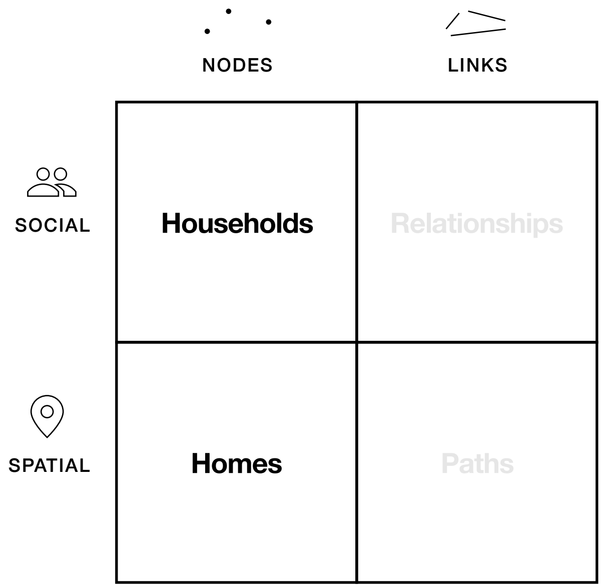

Figure 4.2: Households and homes are the social and spatial map elements that CHWs collected.

4.3 *Effect of App Design on Participants*

Overall, the analysis shows notable effects of CHWs using the app Yego. Its design with list and map representations of the household data enabled CHWs to develop new spatial insights of their communities. For this analysis, we interviewed the participants at the end of the study to gain overall insight on interactivity with the app and how the effects of the design. We used the questionnaire in appendix A and a local translator to conduct the in-person interviews in February 2017 at the end of the field pilot. The quotations are transcriptions of the recorded translations in English, and all participants gave written consent to be directly quoted. The qualitative evaluation of the responses complements the quantitative and qualitative before-and-after analysis in section 4.2 of how the system performed for the CHWs.

Most of the CHWs showed Yego to the members of households that were home when they visited or to their significant others. People were generally curious about the smartphone and were fascinated by the ability to display large images on the screen. After using Yego, all CHWs were able to relate to the map view. The main effect came from seeing the distribution of their catchments' households throughout the village once the list was completely mapped. One CHW commented:

> It's interesting to see how all households in my catchment are distributed along a line. I'm surprised about where I'm located in relation to them. (CHW 1.3, *Field Interview* 2017)

Other CHWs even reported developing the insight that they "could change in what order to visit the households" (CHW 3.1, *Field Interview* 2017). CHWs know the locations of all the houses they serve from memory, but they demonstrated a change in how they think about their catchment. One CHW responded to the question "Did you learn anything new?" with:

> I knew where my households are, but the map changed the representation in my mind. Before it was easy to say "I will first go to this house and next to this one," but the representation of the households' distribution was not as clear as it is today. (CHW 1.1, *Field Interview* 2017)

The app also enabled the CHWs to see their villages in relation to surrounding landscapes and by further zooming out their bordering country, Uganda. Seeing their avatar move in real time within the map helped with situating themselves in the greater context. In general, the map in Yego gave many of the CHWs new insights on how to more efficiently visit their households and provided the first exposure to a top-down view of their villages. One CHW described their experience of seeing themselves on a detailed map:

> Seeing a map of me so close-up is new. Usually, I just see Africa. (CHW 4.1, *Field Interview* 2017)

We asked the CHWs if there were anything else they would like to see or do with the app. Many mentioned that they would like to use the map to locate where important structures were located, including familiar locations like the hospital, the sector and cell office locations, and other villages. Many suggested that it would be useful if the app could give them directions to unfamiliar places. Many were also curious about where other people were located in relation to their map. "I would like to see where my friends are located." (CHW 3.3, *Field Interview* 2017) Another wanted to be able to see the location of other CHWs, and even the location of nurses. One CHW wanted to use the app as a teaching tool, "[I would like to] use the map to explain to children or other people how long it takes to get to certain places and help direct them." (CHW 1.2, *Field Interview* 2017)

In general, the CHWs believed the app had the potential to make their jobs easier. The changed perception of where households were located within their catchment could be used to help plan visits based on priorities and location clusters. Because the app logged the households that the CHW had already visited, one CHW commented that knowing which household they last visited would help with their workflow. Also, CHWs can see which households they have not yet visited and plan their days accordingly.

One of the time-consuming parts of their job is documentation through writing, thus, one CHW suggests that using Yego could "help save time because it is automatic. I could avoid writing which is time-consuming. It would be good to have more tools like [the malaria reporting system] RapidSMS to save time." (CHW 3.2, *Field Interview* 2017) If health information could be integrated into the app, one CHW suggested that "it could be helpful to see on the map which households do not have good hygiene practices so I can teach them." (CHW 4.1, *Field Interview* 2017) If the app were able to show where the other CHWs were located, one CHW would use the app to "exchange news with my peers" (CHW 4.1, *Field Interview* 2017).

Even though using the smartphone's camera was not part of the pilot activity, all twelve CHWs used the camera. On average CHWs took around 30 photos each which they proudly shared with us. Common themes were their homes, families, and self-portraits. An intersting finding was a photo of another CHW's device that was displaying a previously captured photo. Another interesting use case was the video recording of a small TV playing a music video at a commercial center to bring entertainment to the CHW's home. Both instances are forms of physical peer-to-peer sharing of digital data. In section 5.6 we discuss future directions for these observations.

In summary, the design of Yego with its list and map representations of the household data enabled CHWs to develop new spatial insights of their communities that are helpful for their job. The app's ease of use and ability to show a broader view of the surroundings motivated CHWs to map their households and to imagine new use cases. The participating CHWs received direct value by using Yego and see potential in continued use.

4.4 Spatial and Social Analytics

With the foundations for both spatial and social maps established, we examine the data through different lenses to gain a better understanding of the community structure. First, we analyze spatial and social data individually and then bring the two together into a combined network and map.

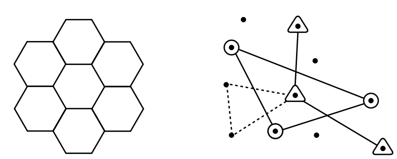

Geospatial Analysis

With the spatial nodes – the mapped house locations – we perform geospatial analysis[5] at a sub-village level. For example, by overlaying a hexagonal grid[6] with a cell size of merely 250 meters on the ground, we compute and visualize fine-resolution population density across the geographic area seen in fig. 4.3.

To understand patterns of life through a spatial lens, resolution matters greatly. Before the CHWs collecting this data, the best pop-

[5] This geospatial analysis is accomplished programmatically using the JavaScript library Turf.

[6] Due to their circularity, hexagons reduce sampling bias compared to a square fishnet grid. Hexagons are the shape with the lowest perimeter-area ratio that still tessellate into a continuous grid.

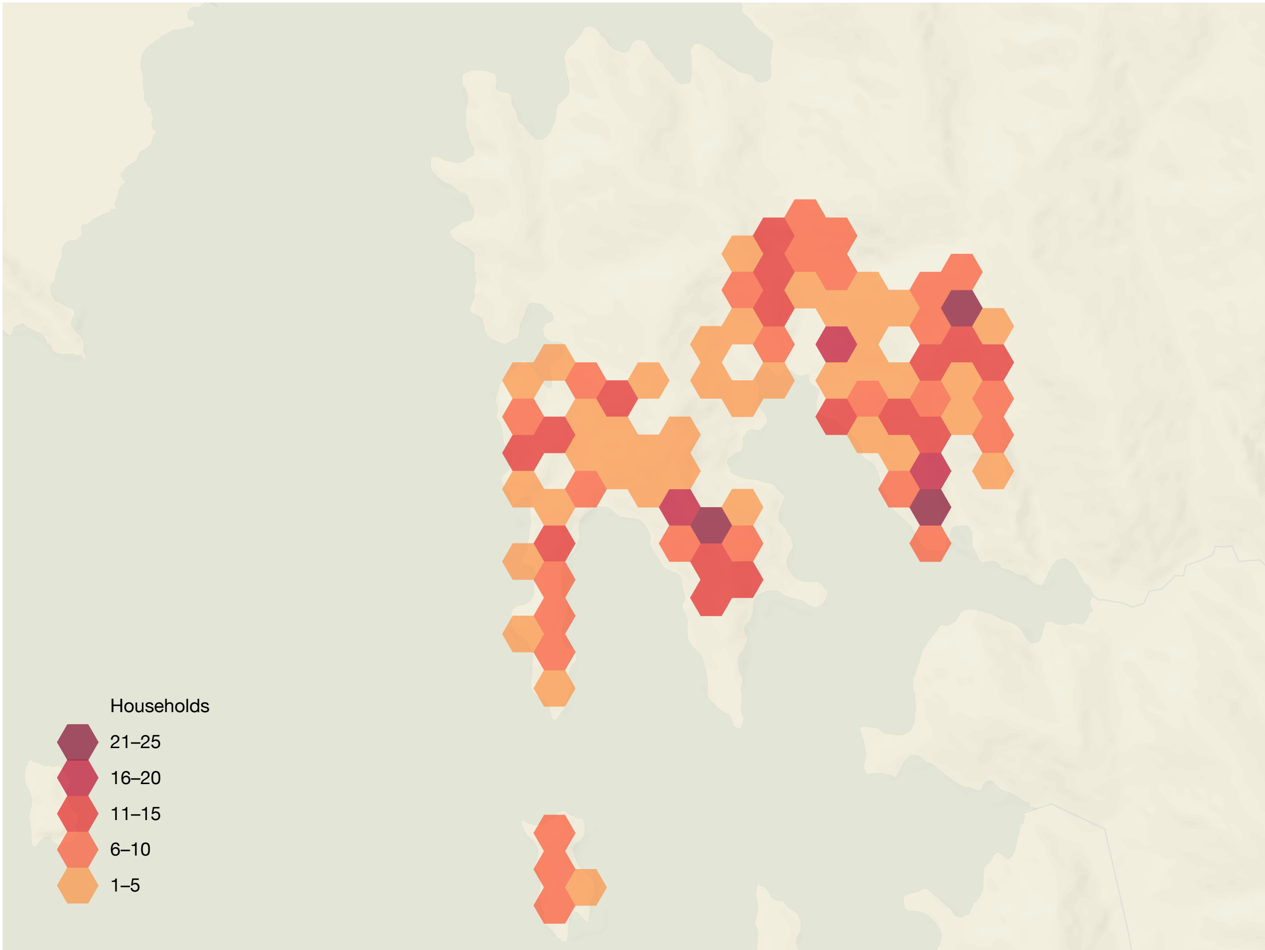

Figure 4.3: Population density map with a cell size of 250 meters.

ulation density representation of this area would merely yield four large, colored shapes (one per village) with similar hues (similar populations). Such coarse and irregularly shaped areas determined by the village boundaries make spatial analysis of other datasets difficult. A fine, regularly shaped grid can be used for hotspot analysis at the household-level, for example, where the focus is on presence or absence of health incidents.

If we bring the CHWs back onto the map and calculate concave hulls around the households that each CHW serves, we derive their spatial catchment areas. This data has not previously been available in a structured way. As visualized in fig. 4.4, we can make spatial patterns of the community health network visible.

When viewing the catchment areas through such a lens, we can make certain qualitative observations. The CHWs are depicted as larger circles, and their catchment has the same color. At first glance,

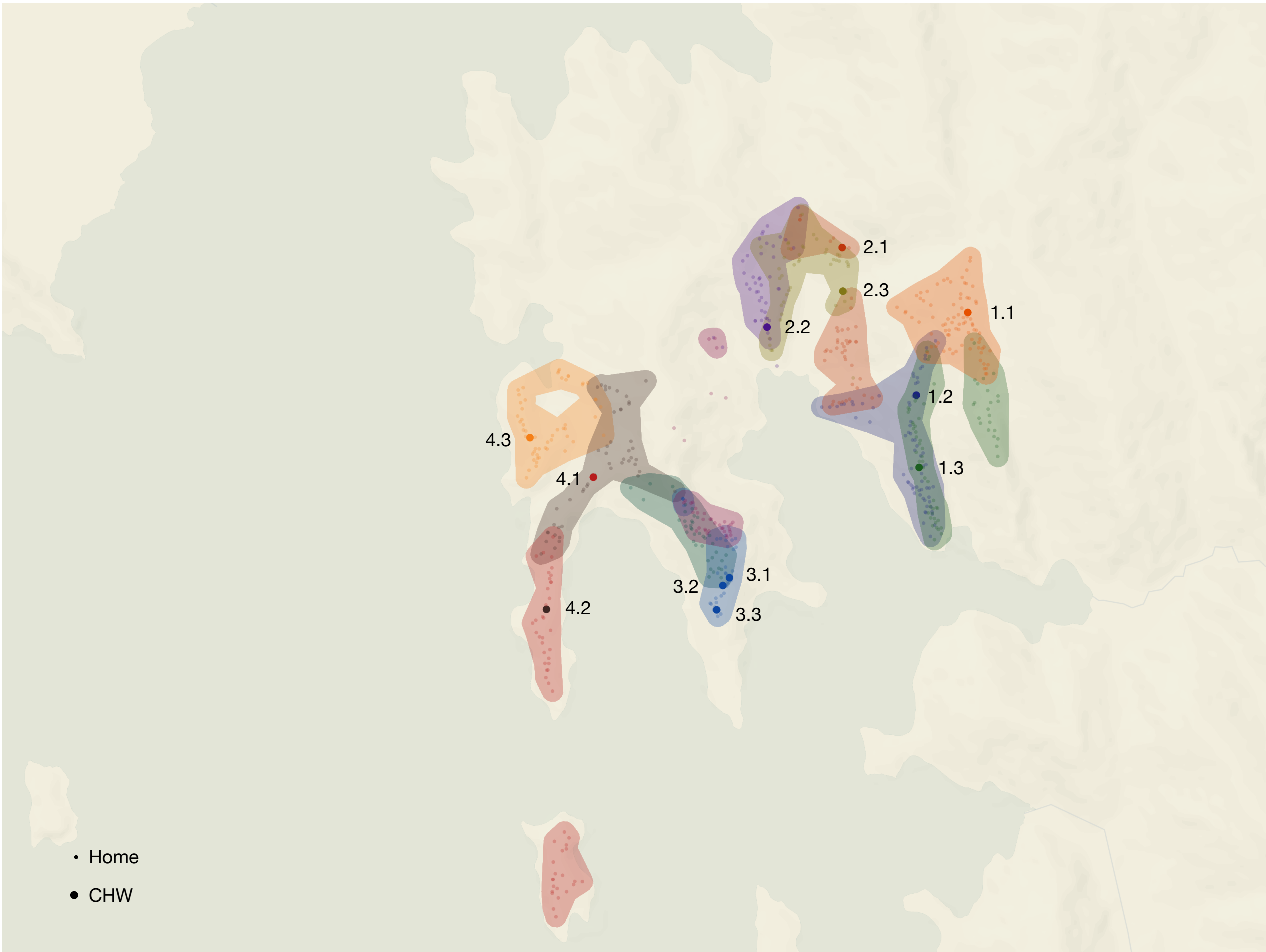

Figure 4.4: Map visualization showing the twelve catchment areas served by CHWs.

it is noticeable that there are overlaps of catchments. Most overlaps occur on the edges, but for others, the majority of the catchment is overlapping with another one. Some catchments are discontinuous – they are split into two parts due to gaps. For example, a significant share of one catchment (22 households) is located on an island only accessible by boat. Multiple CHWs live outside of their catchment. Some of them serve their household; others are served by the CHW responsible for that area (e. g. the three CHWs sharing the same color). Two CHWs live in the center of one another's catchment. Finally, apart from the cluster of three (depicted in blue), CHWs are evenly distributed across the area. These qualitative observations are further discussed in section 5.6 to put them in context and show potential applications.

More sophisticated spatial analysis requires a path and road network.[7] By establishing where humans typically move we avoid treat-

[7] Applications of spatial network analysis include: closest facility, origin-destination cost matrix, location allocation, fleet vehicle routing, service areas, best route.

ing the space as a continuous surface. Furthermore, geography such as altitude and slope can be incorporated into the model. To fill in the path box of fig. 4.5, we establish a path and road network by comparing GPS tracks from the CHWs with high-resolution satellite imagery and trace the traveled paths.[8] We can assign different impedance levels to each segment by classifying them. The four categories we establish are: main dirt road, side dirt road, footpath, and boat (fig. 4.6).

[8] The geospatial network analysis is accomplished using the ArcGIS software and the Network Analyst module.

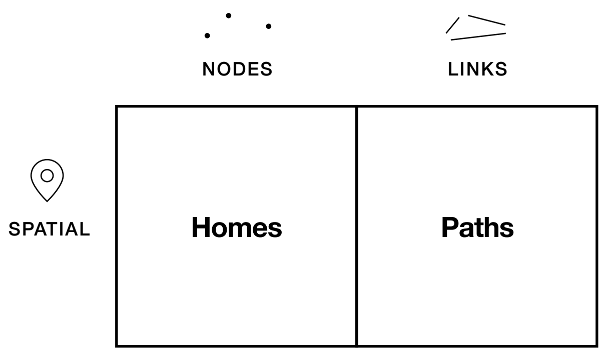

Figure 4.5: Fully mapped features along the spatial dimension.

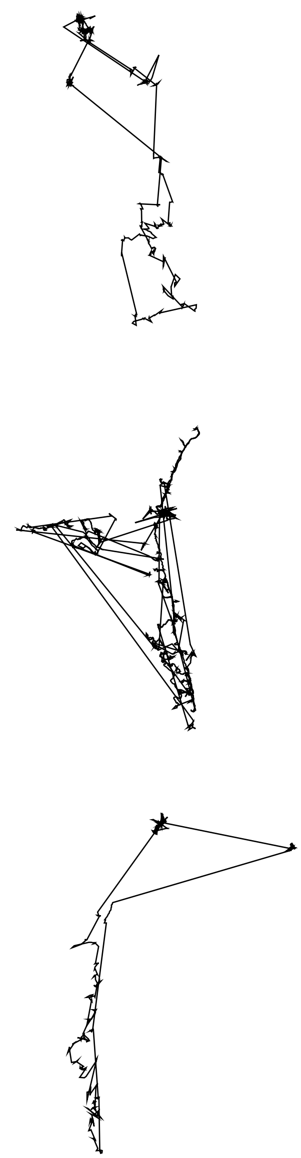

We explore the scenario of spatially optimizing the location of CHWs to reduce the total "cost" to serve the community. This approach is known as the *closest facility* problem. We want to keep the number of CHWs per village at three and calculate their optimal location. We turn all households into potential candidates and likewise all households into demand points. The home locations snap to the closest point on the path network, and all calculations are constraint to the network.

Consider the village "4," Gisirwe. The three CHWs serve 55, 51, and 55 households to a total of 161 (table 4.1). However, the effort to serve these catchment areas might differ greatly. The model nominates the locations of households "53", "49", and "40" to serve catchments with the sizes 80, 49, and 22. As shown on the map in fig. 3.8, the village includes an island. In our catchment map in fig. 4.4 we see that CHW "4.1" currently has to travel through the catchment from CHW "4.2" and serve the 22 households on the island by boat. The model spatially optimizes and places a CHW there to serve the island population directly. This scenario demonstrates the use of sophisticated spatial analytics tools to model suggestions that might increase the efficiency and effectiveness of community programs.

Social Network Analysis

The mapped households are the basis to create the second lens to look at the structure of the community, the social lens. We estab-

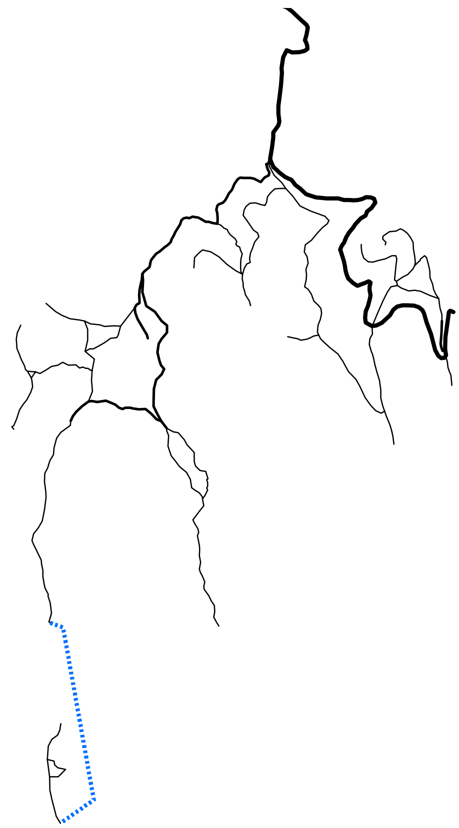

Figure 4.6: A structured network of common human paths enables sophisticated spatial analysis and optimization.

lish links between the nodes to achieve a clearer view. However, the CHWs did not collect any social network data during our pilot. To fill in the relationship box of fig. 4.7, we build a model using a dataset PIH shared with us that contains publicly known associations of households with local institutions. For example, local institutions associated with households include economic cooperatives such as fishing, road construction, or sweater knitting.

The dataset was collected with pen and paper questionnaires by surveyors over the course of approximately one month. Only "group memberships" that are non-sensitive and publicly known (i. e. daily habits) are included. We consider group memberships along seven dimensions: economic cooperatives, traditional ambulance groups,[9] frequented markets, frequented shops, parents of children attending schools, drinking water sources, and cooking demonstrations.[10] With these group memberships, we model the *probability* of relationships between two households[11] – in other words, how likely it is that members of two households know one another.

To join this secondary dataset with our primary data – the nodes collected by CHWs – we first clean and transform the social link data.

[9] A community ambulance *ingobyi* is a traditional stretcher to carry the sick to a clinic. There are about 30 stretchers in the Kaganda cell kept at specific households forming groups of about 50 members each.

[10] Cooking demonstrations are held in each village once a month to educate mothers of children under the age of five about nutritional values.

[11] In contrast, the *strength* of relationships between two households (i. e. how socially close members of two households are) could only be determined by asking personal questions (e. g. names of relatives and friends, or people from whom the respondent gets advice) like Banerjee et al. (2013) did.

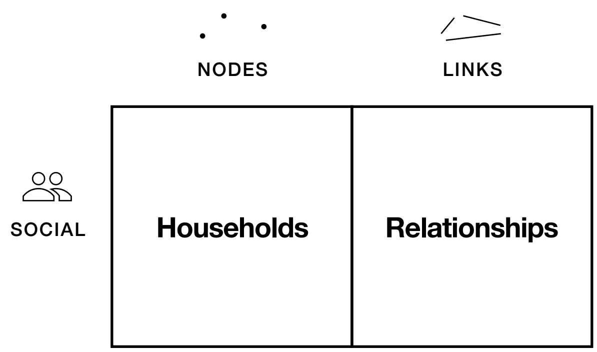

Figure 4.7: Fully mapped features along the social dimension.

Although the data is available as digital spreadsheet, it was entered manually from paper surveys and contains duplicate entries, swapped fields, typos, and varying notations. To standardize fields, we define canonical forms and convert all possible representations into this standard form. For example, we define the canonical form of a repeatedly mentioned group for cooking demonstrations as:

```
NURSERY SCHOOL (MUMPINGA)
```

We then normalize by manually clustering likely representations of that group:

```
MUMPINGA, NURSERY SCHOOL
KUMPINGA, NURSERY SCHOOL
Mu mpinga Nursery School
Mumpinga nursery school
Mu mpinga nursery
```

Next, we build the social network. Using the cleaned social link data, we transform the set of unconnected nodes into a mathematical graph with edges. When creating edges, we distinguish between *direct social links* and *institutional links*. We treat them differently when determining their weight – the probability for a relationship. Direct social links are observed ground truth links (e. g. all twelve CHWs know one another). Their "probability" is thus 100 percent. Institutional links (e. g. two households are members of the same economic cooperative group) are unobserved, and we infer their probability in a later step. All edges are undirected, meaning there is no directedness in the relationship between two households.

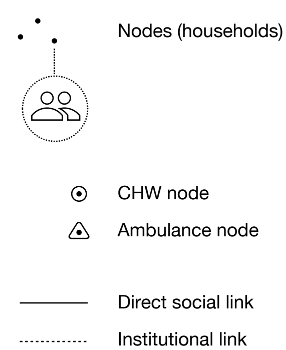

Our social network model differs from location-based social network models such as the one Pelechrinis and Krishnamurthy (2016) describe in that (1) we include direct social links, (2) we form institutional links around spatial institutions (e. g. water sources) and social institutions (e. g. cooperatives), (3) we ground-truth the weight of different types of institutional links using survey answers from CHWs (fig. 4.8).

In an additional module of our survey during the pilot, we asked CHWs[12] on a 5-point Likert scale how often two households, in general, would interact with one another if they are common members of the groups found in the social link data (cf. questionnaire in appendix A). For example, households that frequent the same shops have a higher probability of interacting with one another (mean of answers was $\mu = 4.3$ out of 5) than households that attend the same cooking demonstration (mean of answers was $\mu = 3.0$ out of 5).

[12] Nine out of twelve CHWs completed this module due to time constraints.

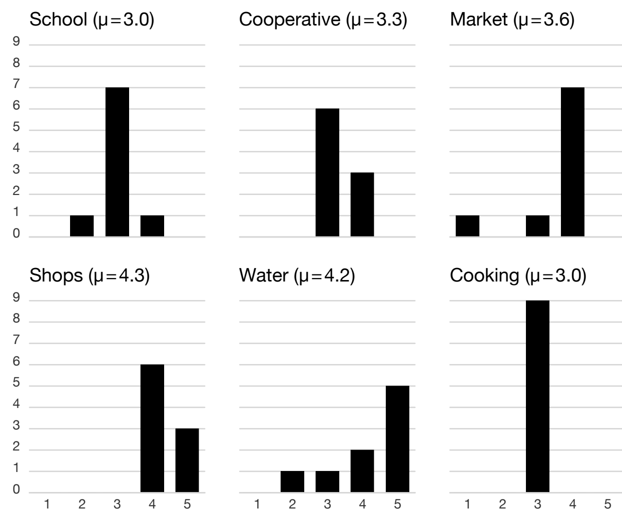

Figure 4.8: Survey answers from CHWs $(N = 9)$ to ground-truth how the social network model weights group relationships.

To initialize the graph, let all 632 households be its nodes.

Step 1 Connect all CHWs (direct social link). CHWs know one another from regional trainings. This creates $n(n-1) \div 2$ edges. For $n = 12$ this results in 66 edges.

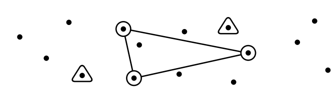

Step 2 Connect CHWs to their catchments (direct social link). CHWs know each household assigned to them from routine visits. This creates number of households minus number of CHWs edges. For $households = 632$ and $chws = 12$ this results in 620 edges.

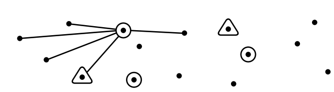

Step 3 Connect members of ambulances (direct social link). Groups meet regularly at the home where the stretcher is kept. This creates $n(n-1) \div 2$ edges for each ambulance. For 18 ambulances with $n = (52, 28, 15, 30, 30, 15, 34, 29, 23, 33, 39, 34, 56, 24, 22, 54, 28, 55)$ this results in 11 175 edges. All social links from steps 1 to 3 result in a total of 11 861 edges.

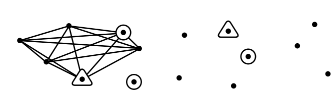

Step 4 Connect all members of the same institutes according to their weights seen in fig. 4.8 (institutional link). Households can be members of multiple institutes and multiple groups per institutional type. Example cooking demonstrations creates $n(n-1) \div 2$ edges for each demonstration. For four demonstrations with $n = (105, 69, 62, 44)$ results in 10 643 edges. All institutional links from this step result in 416 377 edges.

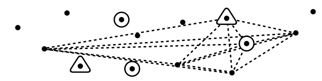

Step 5 Combine all social and institutional links between any two nodes into a single edge with a probability for a relationship. This step combines the original total of 428 238 edges to 179 582 edges (a fully connected graph for $n = 632$ would have 199 396 edges).

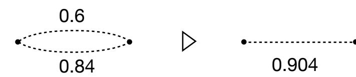

In combining the links,[13] we can discard institutional links between two nodes that are connected through a social link because we know that they know one another. If two nodes are connected through institutional links only, we must make a choice how to use the weights to infer the probability of a social link. We combine the institutional weights into a single probability for a relationship denoted as $\mathbb{P}_r$ by assuming multiple institutional links between two nodes as independent events ("or").[14]

[13] We create a multigraph that allows multiple edges between two nodes. We use Python and the library NetworkX.

[14] A more advanced inference model could account for potential correlations through training on a set of ground truth social links and known non-links.

$$\mathbb{P}_r = (School \cup Cooperative \cup Market \cup Shops \cup Water \cup Cooking)$$

To quantify the results, we assume that if all CHWs responded "5" for all six institutional types shown in fig. 4.8, and two nodes were connected through all six institutional links, the probability be equal to a ground truth social link and define $social = 5$.[15] Example calculation for School ($\mu = 3.0$) and Water ($\mu = 4.2$) using De Morgan's law:[16]

$$\begin{aligned}
\mathbb{P}_r &= (School \cup Water) \\
\overline{\mathbb{P}_r} &= \overline{(School \cup Water)} \\
\overline{\mathbb{P}_r} &= \overline{School} \cap \overline{Water} \\
\overline{\mathbb{P}_r} &= \overline{School} \times \overline{Water} \\
\overline{\mathbb{P}_r} &= (1 - \frac{School}{social}) \times (1 - \frac{Water}{social}) \\
\mathbb{P}_r &= 1 - (1 - \frac{School}{social}) \times (1 - \frac{Water}{social}) \\
\mathbb{P}_r &= 1 - (1 - \frac{3}{5}) \times (1 - \frac{4.2}{5}) = 0.904
\end{aligned}$$

[15] Because for no institutional type $\mu = 5$, two nodes that are only connected through institutional links always have a $\mathbb{P}_r < 1$. We could discount the probability of knowing one another through institutional links by choosing a social link value above 5. To calculate, assume $\mathbb{P}_r < 1$ for full institutional connection and solve for $social$: $\mathbb{P}_r = 1 - (1 - (5 \div social))^6$.

[16] By moving from "or" to "and" we can multiply the weights to get a single probability between 0 and 1.

The probabilities of our edges shown in the histogram of fig. 4.9 follow a multimodal distribution. We can distinguish our likelihoods for a relationship into five categories: "unknown" at 0, "neutral" (neither likely nor unlikely) around the small spike at 0.6, "somewhat likely" around the big spike at 0.7, "very likely" above 0.8 and

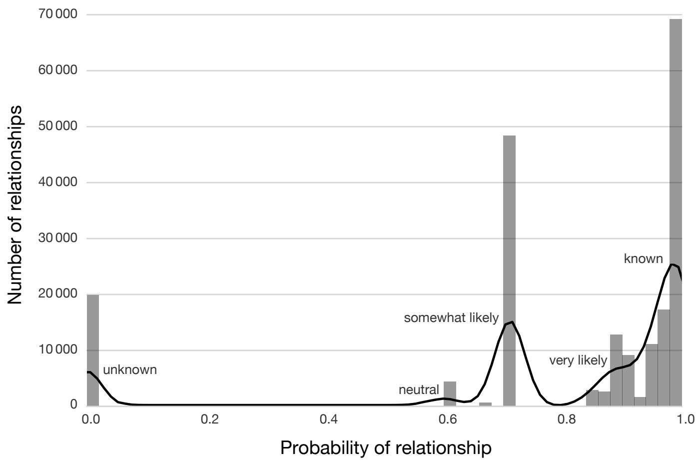

Figure 4.9: Histogram of social link edge weights with overlaid kernel density estimation line to examine the distribution of relationship probabilities.

"known" at 1. Setting aside the categories "unknown" (no link) and "known" (social link), we see a trimodal distribution.

To finalize the graph, we filter the edges to only high-quality edges. We set out to discard "neutral" and "somewhat likely" edges and only keep "very likely" institutional links and all known social links. We examine the percentage of edges to keep in the complementary cumulative distribution chart shown in fig. 4.10. The complementary cumulative distribution shows how often the random variable – $\mathbb{P}_r$ in our case – is above a particular level. The vertical line depicts the numeric threshold that can be thought of as "moving horizontally" to filter the edges. The corresponding value on the y-axis is the percentage of edges to keep. We choose the threshold to keep edges with a likelihood for a relationship above 75 percent ($\mathbb{P}_r > 0.75$). Figure 4.10 highlights the first cluster of probabilities to the right of the threshold mapping to approximately 60 percent on the y-axis. Thus, for the next set of analyses our model considers 126 391 ($\sim$60 % of total) high-quality edges.

With the established social network, we can examine its structure. To quantitatively analyze the importance of nodes in the network we calculate node centrality measures. Different centralities measure different kinds of importances. Common centrality measures that apply to social networks are:

Betweenness Nodes that act as connecting bridges for other nodes.

Degree Nodes that have a high number of connections.

Eigenvector Nodes that are connected to other well-connected nodes.

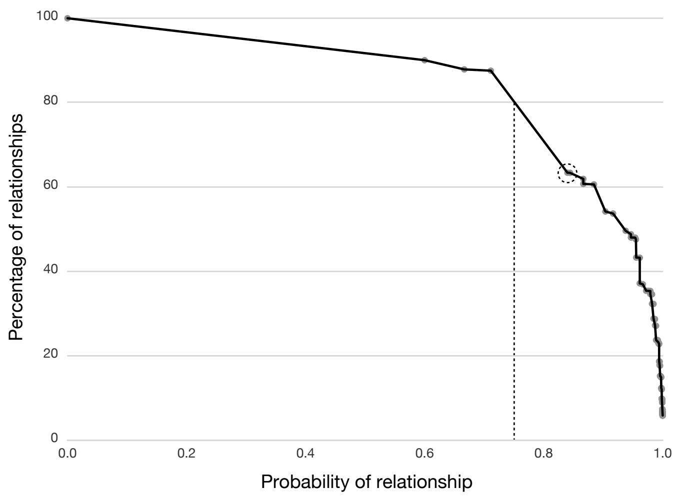

Figure 4.10: The complementary cumulative distribution of social link edges shows the filtered percentage ($\sim 60\,\%$) of high-quality social link edges ($\mathbb{P}_r > 0.75$) to consider in the model.

We calculate betweenness as described in (Brandes 2001) and PageRank[17] (with $damping = 0.85$ and $tolerance = 0.001$) as described in (Brin and Page 1998). Degree centrality is calculated by counting the number of edges connected to a given node. By ranking the nodes based on their centrality value, we identify the most important nodes. Table 4.3 shows the top 30 nodes (top five percent) in the network for each of the centrality measures.

[17] PageRank is a variant of eigenvector centrality and used by Google to rank search results. The name PageRank comes from one of the company's co-founders, Larry Page.

Top nodes differ between the three centrality measures. Many overall high-ranking nodes are from villages "1" and "3." Few from village "4" rank high in degree centrality and the highest nodes from village "2" rank in the 31 to 50 range (not shown in table 4.3). CHWs and ambulance leaders are well represented among the high-ranking nodes.

The number one ranking household (first row in table 4.3) for betweenness and PageRank centralities is CHW "1.2." The fact that they are not ranking number one for degree centrality can be interpreted as them not being connected to the most households but to other well-connected households.

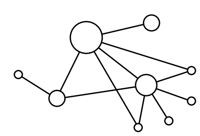

Consider households of village "4," Gisirwe. As shown on the map in fig. 3.8, Gisirwe is the furthest west and located on a peninsula surrounded by water on three sides. For its residents to access shops, schools, or clinics, they must leave their village and pass through other villages. The absence of being on routine walks of other community members is reflected in lower social betweenness centrality. However, households of Gisirwe rank high in degree centrality with "4.3.40" and "4.3.42" even in the top five percent. Presumably, because of the more isolated location, Gisirwe is a tight-knit community.

Table 4.3: Top 30 most central network nodes for graph centrality measures betweenness, degree, and PageRank.

	Centrality		
Rank	Betweenness	Degree	PageRank
1	1.2.1*	1.3.22	1.2.1*
2	1.2.4	3.1.32	3.3.14
3	1.3.20	3.3.14	3.1.32
4	1.3.35	1.3.26	3.2.3
5	1.2.29	3.2.3	3.2.37 †
6	1.2.46	3.2.37 †	3.3.8
7	1.3.22	3.3.8	1.3.22
8	1.3.10	3.3.23	1.3.26
9	1.3.26	1.2.1*	3.3.23
10	1.1.1*	1.2.29	1.3.20
11	1.3.6	1.2.46	1.3.35
12	3.1.32	1.3.20	1.2.4
13	3.3.14	1.3.35	41.2.29
14	1.2.30	1.2.4	1.2.46
15	3.2.3	1.2.30	1.2.30
16	3.2.37 †	3.2.7	1.3.10
17	3.3.8	3.3.18	3.3.28
18	1.3.19	3.3.28	3.2.7
19	3.3.23	3.3.31	3.3.31
20	1.3.43	1.3.10	3.3.18
21	1.3.53	3.2.23	3.2.33
22	1.3.1*	3.2.33	1.3.1*
23	3.2.7	3.3.2	3.2.23
24	3.3.18	1.3.15	3.3.2
25	3.3.28	3.1.30	3.1.2
26	3.3.31	4.3.40	1.1.1*
27	1.1.54	4.3.42	1.3.15
28	3.2.23	3.2.17	1.3.19
29	3.2.33	1.3.1*	3.1.30
30	3.3.2	3.1.2	3.3.34

Node identifier format: `village.chw.household`
* CHW
† Group leader of a community ambulance

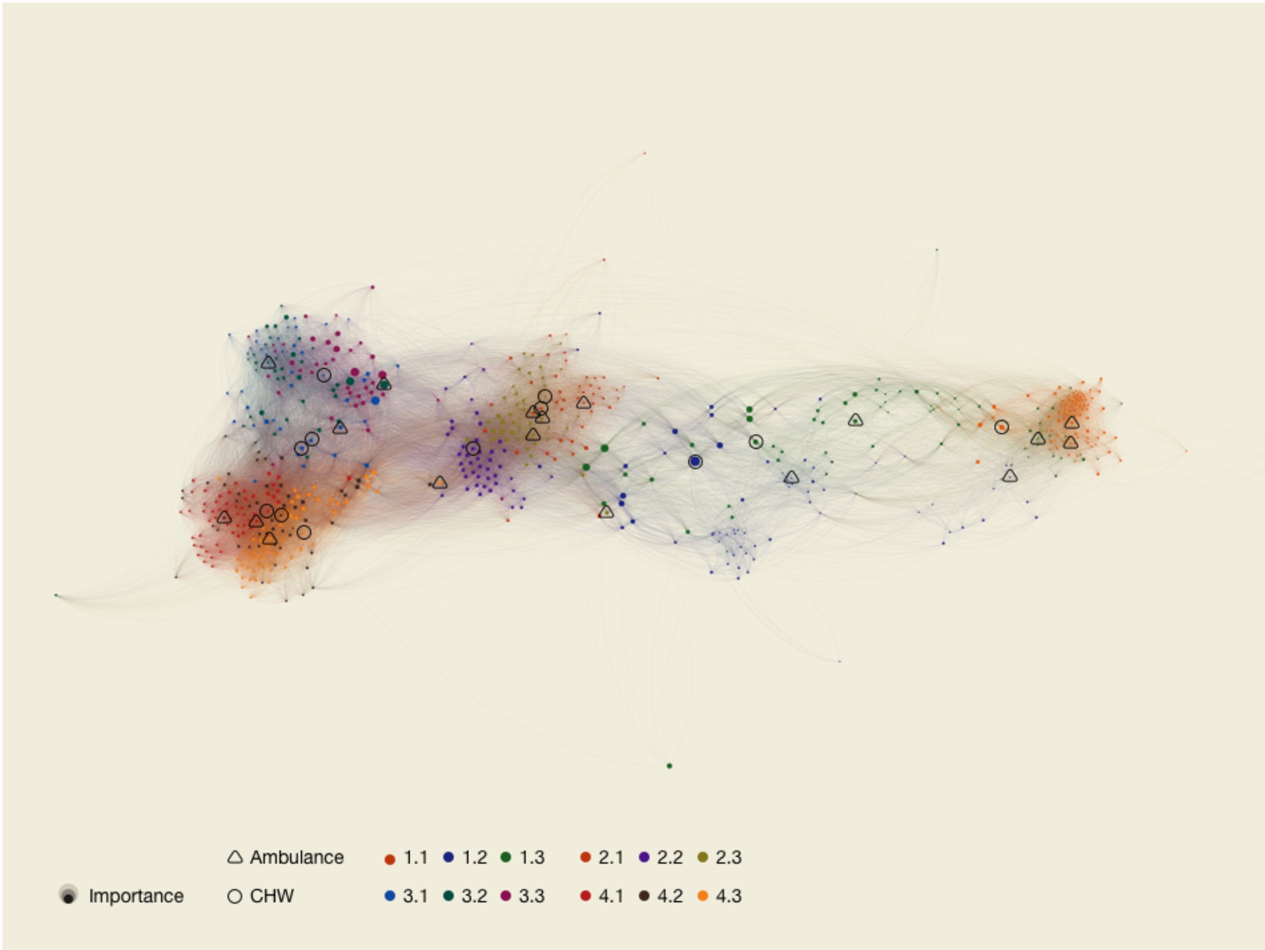

Figure 4.11: Graph visualization of the social community network (632 households). The force-directed layout and size of nodes (social PageRank) reveals additional community structures.

To advance such qualitative interpretation, we visualize[18] the network. The nodes are placed based on their connections in the network. Using a force-directed layout inspired by principles of physics, we create an intuitive lens onto our social network shown in fig. 4.11. Note that there is no GPS data in this model.

Each circle represents a household and each line a social link. The households are colored by the CHW that serves them. The size of each circle is determined by its PageRank. We choose this centrality measure to highlight well-connected and thus influential nodes. As we show in table 4.3 and now represent visually, households with special functions (CHW and ambulance group leader) are influential but not dominating. There are multiple big circles with no special function in the original dataset – in other words, the model surfaces highly connected residents who are neither CHWs nor ambulance group leaders.

[18] The graph visualizations are created with the application Gephi based on output from our network-generating Python script.

CHWs are well embedded in their catchment areas, and CHWs of the same village are mostly located close to one another. Rough village-based clusters emerge based on the social connections. We see an expected correlation between spatial proximity and probability of knowing someone. Other social structures were not apparent from the spatial and CHW catchment map alone.

We notice high social densities for households of CHWs "1.1," "2.2," and all of the village "4;" and low densities for CHWs "1.2" and "1.3." The households of CHWs "1.1" and "2.2" are socially homogeneous while they are heterogeneous (i. e. mixed with households from other CHWs) for all others, in particular, "1.2," "1.3," and all of the village "3." Another structure that this social lens reveals is the likelihood of entire groups knowing one another – the further apart, the less likely. The most noticeable distance is between the tight-knit cluster of CHW "1.1" and the rest of the community.

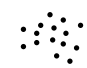

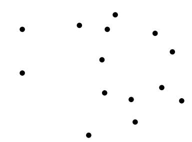

A Socio-Spatial Community Network

To further identify and compare community patterns, we ground the social model in spatial reality. In retaining all attributes of our social map – the node's importance, color, and connections – but pinning the nodes to their spatial coordinates collected by CHWs, we create the socio-spatial map in fig. 4.12. We generalize this notion of a socio-spatial map in section 5.1 as a map that combines people, places, paths, and relationships. Such a map enables deeper analysis of human networks in their physical environment and expands the definition of being visible.

The map exhibits nuances and exceptions to the correlation between spatial and social centrality. For example, although CHW "1.1" lives spatially in the center of their catchment area, socially they are located closer to the rest of the community (cf. fig. 4.11), acting as a bridging node for their catchment. The map further reveals that this socially homogeneous cluster is a big group (82 households) living separated by the main road from the rest of village "1" and at higher altitude up on a hill.

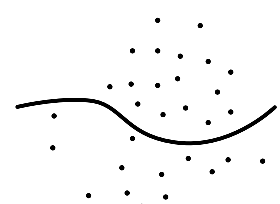

While some important households are located near centers with commercial activity (left in fig. 4.13), other socially central households are spatially not central in villages (middle in fig. 4.13). Consider the households in the southeast of the big map at the shore. Their high social centrality is likely due to many households fetching water in that meadow daily (right in fig. 4.13) – a dimension that was captured in the social link data.

This simple model[19] predicts patterns that align with our empirical evidence, and offers insights into community structure that

[19] Without any additional data, future models could factor in group sizes (favor smaller), dated events (decay over time), or triadic closure across households and groups (i. e. ties between A–B and A–C increase probability of a tie between B–C).

Figure 4.12: Socio-spatial map visualization of the community network (632 households). The social community structure is mapped onto the spatial environment.

would be difficult to see on the ground and at scale. We would like to acknowledge that many dimensions of reality are not included in such models. Ideas how these socio-spatial observations can be used, though, are discussed throughout section 5.6. For example, by creating a community-based interface to confirm social links, we could spin the human–machine loop one more time and let the community validate and influence the model.

Figure 4.13: Detailed views of community centrality: a commercial center (left), a socially central household outside the village center (middle), and high centralities close to a water fetching location (right).

5
Discussion

We discuss the nature of social and spatial maps and how they expand the notion of visibility. We generalize the diagram of map elements to people, places, paths, and relationships to define a socio-spatial map. Furthermore, we derive design characteristics of our implemented "social machine" that created this map and new visibility. The derived characteristics such as *Use familiar data to teach unfamiliar representations* center around our human–machine collaboration approach. These characteristics are intended to help others to reason about design decisions in related contexts and to combine them into new systems. Finally, we discuss sharing of the novel map data with different parties, the limitations and scaling potential of our system, and the potential impact and future directions of this work at the intersection of social and spatial mapping.

5.1 Socio-Spatial Maps

The maps created from this thesis work – spatial, social, and combination into socio-spatial – provide different lenses to view a previously unmapped region. In section 1.1 we define digital invisibility and the impact of becoming visible for these communities. The combination of geographic and social visibility into a multimodal map expands the definition of being visible. In addition to the ground-validated maps and their analyses, this thesis proposes a framework to reason about what we call socio-spatial maps. In this section, we discuss what a socio-spatial map defines and how it enables us to look at maps differentely.

Through human–machine crowdmapping we mapped homes and households; through data analysis, we mapped paths and relationships. Homes and households, paths and relationships, complete the four quadrants of our 2×2-matrix-diagram. The diagram introduced in section 1.4 generalizes to any spatial map replacing homes with the notion of *places*. Analogously, the diagram generalizes to any social map replacing households with the notion of *people*. Figure 5.1 shows these four map elements: a socio-spatial map consists of people, places, paths, and relationships. Combining social and spatial into one diagram allows us to navigate between pairs of socio-spatial map elements and examine the nature of their combination. We can see how the different elements relate in the abstract and discuss examples of specific maps.

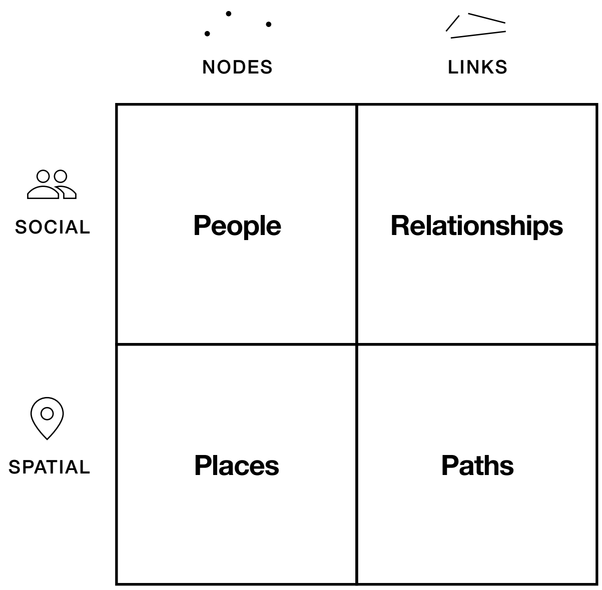

Figure 5.1: The socio-spatial map diagram combines social and spatial maps. A socio-spatial map consists of people, places, paths, and relationships.

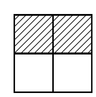

People & Relationships. Any human network consists of people and their relations. Relationships can have different meanings in different social maps, but they always show how people relate to one another. The different meanings can be social proximity, the frequency of communication, or, in the case of our map, the likelihood of knowing one another. Some meanings can be quantified while others cannot. In a view of the world that captures the physical, relationships between people remain the one invisible element. We can only infer relationships, or observe their effects, such as trails left behind by people traveling to meet other people. People digitally self-mapping relationships lead to vibrant present-day social networks, demonstrating an opportunity for any community to adopt such "maps" by defining relationship meanings relevant to them. For example, in a rural setting at the BOP with little formal education, it is valuable to share knowledge. Subsistence farmers could peer-to-peer share agricultural best practices or live soil data. Another example is maps showing general human contact for epidemiology. Contact tracing could help CHWs to diagnose their patients. Maps

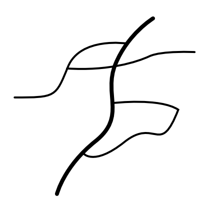

including both elements along the social dimension are among the most valuable and powerful datasets.

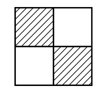

People & Paths. People move along paths and are physically connected with one another through paths. There are paths on land, water, and in the air. Some paths require vehicles to travel them. There are mapped *and* unmapped existing paths, and new paths yet to be explored and mapped. "For many people, [paths] are the predominant elements in their image; [...] along these paths the other environmental elements are arranged and related." (Lynch 1960)[1] Increasingly, people use automatic directions from mapping services to move along mapped paths. People moving along *unmapped* paths – in vehicles or not – offer an opportunity to map these paths automatically. For example, CHWs recording GPS tracks using smartphones could be the basis to map all paths to homes in rural areas economically. Dominant paths would emerge by overlaying the data of multiple people over time.

[1] Lynch (1960) divides the mental image of the city into five elements: path, edge, district, node, landmark.

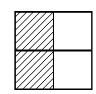

People & Places. The node elements are the primary elements of the socio-spatial map. In some maps, links might carry more information, but links are always a means to interact with another node. If our intention is to navigate the physical space, for example, we are temporarily more interested in the information stored in the map's links to learn where a node element is situated on the greater map. Similarly, if we want to learn how people are situated in a greater community, we examine the map's links. In creating our socio-spatial map, each collected data point mapped a place *and* the people living in it, establishing the nodes along both dimensions at once. However, this relationship we leverage between social and spatial nodes is not fixed: if someone moves away, new people might move into the existing place. If a structure is removed, one must still account for the inhabiting household. Unlike paths that are commonly shared, the majority of places are privately owned by people. The tradition of naming public places after people is digitally mirrored in products such as Foursquare that hosts leaderboards of people's physical "check-ins" at places.[2] Considering the interaction between people and places – meeting someone at a place – we can turn either social or spatial node into the variable: Uber's "People are the new places" turns a person into a destination; visiting a fixed place to meet *anyone* does the inverse.

[2] Typical check-ins capture presence at a place and a time. Expanding this notion to a "map as memory store" could provide an authentically mobile spatial plane for anything (cf. how the web was designed for documents and desktop computers).

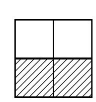

Places & Paths. The classic cartographic map combines both elements along the spatial dimension. Places and paths are the most visible, and hence their mere existence has the biggest potential to be mapped with images from above, as we do using satellite imagery and machine learning. Places can be small or large and can have a two-dimensional extent or just be a point feature. Lynch (1960) distinguishes between places we can enter (districts, nodes) and places we cannot (landmarks). Some places are the result of junctions or other convergences of paths. The length of paths determines the proximity of places. Depending on the point of view and mode of travel, the proximity changes (e. g. uphill compared to downhill). Paths around a significant place commonly borrow its name, whereas most places[3] are addressed by the name of the path leading to them. Some paths in big places like cities are named after faraway places *to* which they lead.[4]

[3] Addresses in Japan do not reference street names; they start with the largest geographical entity and proceed to the most specific one (typically a house number).

[4] For example, the primary street leaving Bern, Switzerland towards the city of Thun is called *Thunstrasse*. The name changes midway, and in Thun, the same street is called *Bernstrasse*.

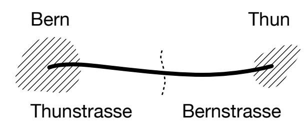

Places & Relationships. Relationships are often formed around physical places. To build our community social network, we use public institutions – some tied to a place and others not – to build affiliations for groups of people. Location check-in data could facilitate this process from a bottom-up perspective. Known relationships can reveal more about the "nature" of a place (Pelechrinis and Krishnamurthy 2016), such as how likely a place is to facilitate the forming of new relationships. Public places also influence the relationship we have with people living nearby. An example from our socio-spatial community map is the social centrality of people living close to water wells. Inverting the direction between relationships and places, we can ask the question "What is the 'place' of a relationship?" An answer might be a location in the middle of the people at both ends of the relationship or a significant place that brought the relationship into existence.

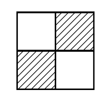

Paths & Relationships. The link elements are the connective tissue of the socio-spatial map. Our map shows that spatial proximity – the length of a path – correlates highly with the probability of a relationship. We assume this is truer for societies in which the main mode of communication is in person and that the length of a path becomes less of a predictor for a relationship when communication happens at a distance. Similarly, the path length is less relevant when we move along paths with vehicles, and especially when flying. Travel and telecommunication "virtualize" relationships and the correlation softens. Commonly, the two domains of social and spatial are treated separately. We can study such patterns at the intersection of social and spatial by combining these two domains and examining them together.

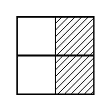

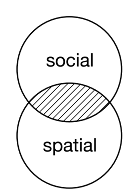

5.2 *Social Machine Characteristics*

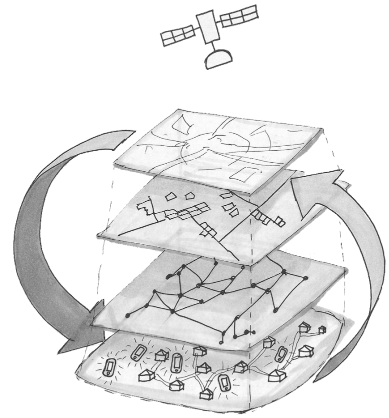

In designing a "social machine" for BOP communities to crowdmap using a human–machine collaboration, we derive a set of its core characteristics in this section. A social machine is a system of humans and machines collaborating to accomplish a task that one part could not accomplish alone (Hendler and Berners-Lee 2010). The design combines the elements of machine-driven analytics and human-powered annotations into a holistic system. Using low-cost smartphones and an intuitive mobile app closes a crucial data and knowledge gap for BOP communities to self-map.

Any social machine is embedded in an existing social fabric. Thus, the technological solution requires a social system to be successful. The effort to design, implement, and nurture such a social system dwarfs the same on the technical side. We identified such a social system – CHWs with their local knowledge and routine household visits – and built our solution on top of it. Doing so is essential for this social machine to work. In acknowledgment of the sophistication of the system's "social" aspect, we now turn to and discuss the "machine," or "human–machine," part that we designed.

In taking this machine apart – from design through implementation to deployment – we derive its core characteristics. This set of core components is intended to help others to reason about design decisions in related contexts and be combined into new systems. We hope to contribute towards removing the barriers that currently create solutions with shortcomings in the domain of digital tools for BOP communities by describing the following set of core characteristics:

START WITH A FAMILIAR CONCEPT. Our mobile app design is accessible to its users because the central interface metaphor is a familiar concept. Although there has never been a smartphone in the communities for which we were designing, we drew from an existing vocabulary. No design exists in isolation. We determine design elements either to be working intuitively, without previous exposure or to reference familiar elements of the community's environment. For example, for buttons, we rely on their visual affordance to push them. To represent homes in the community, instead of a map, we employ a skeuomorphic[5] representation of CHWs' paper-based household lists. In this case, a paper list is represented digitally as a scrollable list that displays the familiar data. The spatial mapping happens implicitly and can be explored in a secondary view.

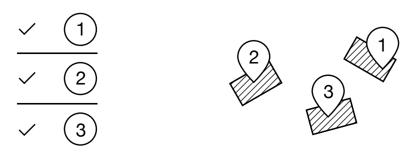

[5] Skeuomorphism is where an interface element mimics its real-world counterpart. Skeuomorphism as we recommend it is different from visual ornamentation.

Express standard concepts with standard representations. When choosing how to represent familiar concepts, the visual language should not differ significantly from established forms. Establishing new forms[6] creates different barriers. We only use basic symbols that are generally understood and have the same meaning across cultures. Although we use little text throughout the interface to account for illiteracy, certain elements remain best expressed textually or numerically. Short labels with numbers, dates, or recognizable, almost "symbolic" words are powerful concepts and are unmatched in their clarity and universality. In today's world, exposure to text is unavoidable, even for the illiterate. By employing text following these guidelines and localizing the content it represents, we build on-ramps for learning and avoid isolation. However, text can be a barrier in other ways than literacy. Translation into low-resource languages is challenging. Hence, an interface that functions without relying on entire sentences scales better to other locales.

[6] For example, representing time with sun or moon icons or drawings. New forms could *support* representation and understanding of a concept.

Use familiar data to teach unfamiliar representations. When we examine a problem from different angles, we use familiar representations to understand unfamiliar data. The inverse works too: moving familiar data into different representations helps us in understanding an unfamiliar representation. We can see how the representation exhibits the data's familiar properties. For example, CHWs are deeply familiar with their environment and the household list but not with cartographic maps. By enabling CHWs to explore households on the map, and providing continuous transitions between the two representations, users can incidentally develop map literacy.

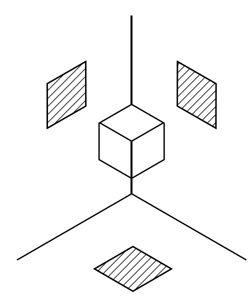

Align use with existing behavior. Technology is generally an amplifier of preexisting human drives. Forming a new habit or adding a new activity to someone's life is challenging. Aligning technology to existing human behavior is the difference between adding a chore and adding an empowering tool. We leverage the existing "human pulse" of CHWs routinely visiting households in their catchment area. Human-centered design is a powerful framework to identify such opportunities. When designing user flows that are companions to existing activities, making steps optional is critical[7] so that a solution can react to real-world situations.

[7] An earlier design was based on families taking self-portraits to initiate GPS location acquisition. With our analyses showing that only 40 percent of heads of household were home during CHWs' visits, the system performance would have suffered.

Teach through imitation, not instruction. Certain skills are better shown than told – learning how to ride a bicycle, for example. With the right protective gear in place, we best learn it by imitating others. The actions to successfully use a smartphone for the first time are easily demonstrated but difficult to describe. For someone unfamiliar with the concept of a touch screen, a manual illustrating its use is difficult to comprehend. Elaborate on-screen instructions are less effective[8] than a brief in-person training. Thus, in addition to aligning the use of technology to existing social behavior, the rollout strategy has to be aligned to existing social structures. This leads to the hypothesis that merely airdropping technology is not effective.

[8] Early intuitions were to put a "teacher" into the phone using localized video messages. Someone *in* the phone teaching the user *about* the phone performed poorly in our user tests.

Focus on a single user action, delegate the rest. It is important to keep the feature set focused and to ask for one user action per feature.[9] For tasks that cannot be accomplished with a single action, a path to delegate subtasks should be available. Delegation can happen to other humans or technology. It can also be getting support by another human who is more knowledgeable. Increasingly, we can delegate much of a task to machines. If technology does the administration, people can focus on doing the creative work (Hendler and Berners-Lee 2010). Our app delegates the recognition of building shapes from satellite imagery to the machine. Freed from the mundane and labor intense task of drawing polygons around structures, users can focus on mapping in the physical world. Secondary interfaces for the human to override the machine can correct mistakes and help the machine behavior improve.

[9] Asking a novice iPhone user to take a photo demonstrates how the camera app violates this principle. The interface grew complex. Ideally, there should be one button: the shutter.

Provide delight. A strong design is not only intuitive but also delightful to use. The design must first be functional, reliable, and usable, but a "usable interface" can be compared to merely "edible food."[10] Usability is a requirement, but not sufficient for a delightful experience. In our case, there is no functional need to show user profiles. The system could work reliably without displaying a map. The app would be usable without animating the users' picture to their actual location. Yet, we incorporate these elements to let CHWs that are engaging in this experience feel empowered and motivated. When looking for elements that provide a delightful experience, the video gaming industry is an obvious inspiration. Providing a delightful experience is the purpose of a game. These techniques tap into core human drives that motivate us, whether in the context of games or not. Hence, such techniques are often described as gamification. Gamifying a software is not just about adding points, badges, and leaderboards to the experience. We employ elements in our UI design (see section 3.4) such as ownership, accomplishment, and creating greater meaning.

[10] "Usable = Edible" is described in Designing for Emotion by Walter (2011), at the time Director of User Experience at MailChimp, which is a great example of bringing joy to productivity software.

Build it sturdy. Consider the context software is used in: user, activity, environment. A common scenario that consumer software is designed for imagines a digitally literate person sitting on a sofa with dedicated time. Such an image has design implications on numerous dimensions. Our scenario is a low-literate person serving patients outside their home in a rural area. In such a context, a good product is primarily a sturdy product. Preparing hardware devices with rubber cases for outdoor use must be followed by optimizing software interfaces for outdoor use. Gaining this mindset provides us with a framework for making design decisions. Specifically, designing with high contrast for high legibility, using large icons and large text. Further, we observe the importance of optimizing the design for one-handed gestures that are visually discoverable and not relying on double-taps nor edge-swipes. Extending hit areas of controls beyond their appearance avoids missed taps with great effect and no visual differences. Unintended actions can be avoided by reducing competing gestures such as scroll *and* tap in a list. A rural farmer's touches on the small glass surface often differ to the gesture detectors' calibration. It is the software's responsibility to account for that. "Sturdy" as we recommend it is not dull, not simplistic. Rather it is simple and reliable.

Be resourceful. A sustainable solution can be designed by decreasing resource consumption over increasing resource production. The limited infrastructural resources in our case are power and connectivity. We avoid draining the battery by engineering the software natively and increasing the time the device's chips operate in low-power mode. We eliminate the need to transmit data to the cloud. Big data powering the app can be preprocessed and bundled with the installation. Other data that is required for the user experience can be acquired using the device's sensors. The camera, GPS, or pedometer, for example, do not require data connectivity to generate new data. The data model and persistence must be architected to allow reliable offline use. If limited connectivity is available, tiny bits of information can be transmitted to signal the existence of larger data that can be transferred physically or by connecting to a central hub.

Enable data export. When a social machine generates data, an important characteristic is for humans to be granted transparent access to view their data. If the data is directly valuable to a stakeholder, exporting directly from the running system should be an option. The system must export the data in an interpretable format. If the stakeholder is unable to receive or interpret digital data altogether, a process to create physical representations is critical in returning the data to its originator. In our case, because CHWs benefit from the mapping data they collect, we allow for direct data export from the smartphones and printing of generated paper maps. Leav-

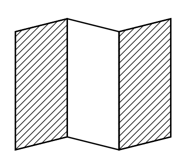

ing these physical maps with the community is the essence of the human engagement and closes the human–machine loop.

5.3 Sharing Map Data and Direct Impact

We shared the map data with the community, with PIH, and are making it accessible to the public. For each party, we considered what an accessible format constituted. This increases the direct impact the map data may have.

For the Community

For most CHWs, using the map view in Yego was the first interaction with any detailed cartographic map representation. Because the app enabled CHWs to transition between familiar and new representations of the collected data, users can incidentally develop map literacy. In addition to learning cartography, we have shown in section 4.3 that by using Yego, CHWs may have gained better spatial understanding of their community. This has direct impact on their CHW activities because it helps to plan household visits based on location clusters more efficiently.

To share these empowering maps with the CHWs permanently, we created physical maps. We leveraged the app's export capabilities as described in section 3.6 to transfer the collected data to a laptop computer. This enabled us to create a printable map representation of the data while still at the CHWs' homes. The map shows the CHW's mapped catchment and their profile to credit their work. One of the twelve maps is included as an example in appendix B. For each CHW we printed their map using a portable printer[11] we brought to the field. Impressions from this process can be seen in fig. 5.2. Ssmallcapshortly after receiving these documents, the community often gathered to find themselves on the map.

[11] In evaluating portable printers, we were looking for a compact color inkjet printer and recommend bringing as many extra ink cartridges as possible (e. g. to print photos).

In addition to sharing the map data with PIH and the public, returning the data to the community in an accessible representation empowers them directly to see and situate themselves on the map. While the institutional value for maps was well understood at the beginning of this study, this process demonstrated the simultaneous value for the individuals to map themselves. The moments when the physical maps emerged created powerful emotions for everybody involved.[12] It appeared that people wanted to find themselves on the map, that they wanted to be visible.

[12] To show our gratitude to the CHWs, we also designed and printed participation certificates for them upon completion of the study.

Figure 5.2: Impressions of the field map printing: a CHWs receiving their map (top), the printer setup (bottom left), and community members studying a map (bottom right).

For Partners in Health

The most detailed map of Burera thus far was made by PIH. The map is displayed on outside walls of public health centers throughout the district. A photo of it can be found in appendix E and the process of creating this village-level map is described in section 2.1.

Upon completion of the pilot, we shared the digital dataset with PIH's GIS expert. Our house-level map has now been integrated into the Burera district administrative map. The updated map included in appendix D shows the pilot area mapped by the CHWs situated in the larger district. Further use of this map is discussed in section 5.6.

For the Public

The entire geospatial dataset will be accessible to the public from June 2017 digitally on the web:

OpenStreetMap.org We are working with OpenStreetMap on an importing strategy to make the dataset part of the mapping service.

VisibleCommunities.MIT.edu We are working on a project website with a section to view and download the dataset in common formats.

5.4 Limitations

Downsides of Mapping

Increased digital mapping of our world – geospatially and along all other dimensions of life – leads to increased visibility by its citizens. This new visibility increases horizons, empowers individuals, and is a powerful tool for organizations. However, making all aspects of life more visible and more transparent can also have downsides. Dennett and Roy (2015) compare the effects of increased transparency in today's digital age with the effects of increased biological variety due to sudden sunlight[13] on life:

> What about the risk of destroying the integrity or effectiveness of organizations by exposing too much of their inner workings to the world? [. . .] A biological perspective helps us see that transparency is a mixed blessing. [. . .] An animal or plant does not have to worry about its cells jumping ship or starting a mutiny; except in the case of cancer, the cells composing multicellular life-forms are docile, obedient slaves. People, in contrast, are individually powerful and intensely curious about the wider world.

[13] The Cambrian Explosion theory is that due to the oceans suddenly turning transparent roughly half a billion years ago, life evolved eyes and as a result variety of different life forms surged.

A completed socio-spatial map creates transparency in multiple dimensions. We took precautionary measures to protect data privacy. This was front of mind for us, especially given the history of abusing registers to identify ethnic groups during Rwanda's 1994 genocide. Fortunately, this particular data is no longer tracked in Rwanda.

Generally, as outsiders to these communities, we only have limited insights into local dynamics (e. g. details of private and governmental land ownership) of places where our system could be deployed. Anecdotally, with our approach – working with and trusting members of the community that are themselves trusted by the larger community – we did not encounter any opposition by community members. However, it is important to consider cases where the increased transparency of mapping could have negative consequences and consider interventions accordingly.

Map Features

Maps are always limited to the perspective of the world that the mapmaker introduces. The ground-validated features of our map are limited to buildings where households live that are served by CHWs. Due to the total coverage of the community health network in that region, we can assume that every residential building is mapped. Because there are almost no structures other than residential buildings in that region, the vast majority of all structures are mapped. The exceptions are a few shops in the two commercial centers in the region. We estimate there to be around a dozen buildings each. Bigger infrastructural buildings like the health clinic and main school in that cell were previously mapped by the government.

Yego does not currently enable ground validation of infrastructure; however, the bundled building footprint detections offer an opportunity to add this feature seamlessly. Our machine learning pipeline detected buildings of all types and this data is bundled with the app. Our design intentionally supports mapping not only from the list view but also from the map view. In this bidirectional interface, tapping a still unannotated building on the map (red footprint) could offer an option to add a new item to the list, or mark it as a mistake (not a building). In addition to the mapped households (green footprint), these non-residential buildings could be colored blue, for example. The difference of the top-down building location predictions and the bottom-up field mapping could estimate unmapped infrastructure.

When mapping additional types of structures, user-generated content could be used to describe them from a ground perspective. Media that is conveniently captured with smartphones in the field could bring a map alive. To map a water well, for example, photo, video, or audio of it could be acquired in the field and attached to that particular location. This has the potential to facilitate a virtual decision-making process for communities to decide what is mapped and how it is represented.

Pilot Scale

Twelve CHWs placing around 3000 people on the map in roughly one week is a meaningful result with a direct impact. If we consider how much beyond the pilot site is left unmapped, however, the scale of our pilot is modest. Our analysis is based on a small number of participants and observations. It is also difficult to estimate the potential impact of a novelty factor or observer effect. Part of the success of our study is attributable to the benefit of having the high accountability structure of the local PIH operations and government in Rwanda; interventions are generally adopted, and data reported reliably. A local scaling-up of the pilot is realistic with the current hardware and CHW operations in the region of Burera. Expanding to all seven villages in Kaganda and mapping the entire cell with its 1045 households requires negligible additional effort. As we elaborate in the next section, the elements of the system have potential to scale further up.

5.5 *Sustainability and Scaling*

A successful field pilot relies on multiple factors: a use case, working technology, and users willing and capable of adopting the solution. This last part – humans using the solution – is the most important factor for adoption, especially for a sustained and scaled-up adoption. Working with an organization that has previously established a social structure to build upon – PIH and the community health network in our case – was critical for the pilot to be successful. When deploying the technology without changes in the longer term and to more communities, we think it will be essential to keep the social aspects (e. g. who to provide the app, how to train them, allowing access to other functions of the device, the duration and nature of the activity, sharing the data) intact to achieve comparable bottom-up enthusiasm and a high-quality outcome. To scale the solution to different regions or countries, we recommend looking for a similar partner that is familiar with the local communities.

The key elements of our system are designed to scale or have the potential to scale in the near future. The human side of the system scales by situating the solution locally and close to the base population. Local capacity is more effective and more economical than relying on outside data collectors or GIS professionals. The technology platform scales atop three trends: (1) rising smartphone adoption in low-income economies (see section 2.2), (2) frequent and economical imaging of the entire earth at a high resolution (see section 3.2), (3) commoditizing high-performance machine learning capabilities (see section 3.3). The reviewed literature suggests strong continuation of these three trends, facilitating a scaling-up of the technology.

In aspiration to start the process of scaling upon completion of this thesis work, we have donated all equipment (twelve complete field kits, cf. section 3.5) to the PIH Rwanda research department. We have familiarized our main collaborator with the technical details of deploying the system and aggregating the data. With this modest equipment, traditional data collection[14] can be outpaced by a factor of ten.

[14] PIH assumes a hired data collector to survey six households per day. At that rate, mapping our 632 households would approximately take 100 person-days of work.

Furthermore, if twelve CHWs put 3000 people on the map in one week, we can extrapolate that they can place up to 156 000 on the map in one year (52 weeks). Even more significantly, because CHWs are present in every village, we can extrapolate that roughly

3000 CHWs can put up to 156 000 people on the map by each mapping their catchment with an average of 52 households in just one week. A "training the trainers" approach can set the required conditions for such a distributed and parallel strategy.

We see it as technically feasible for PIH to scale to more villages, cells, and districts. The potential in Rwanda alone is significant: scaling from four to all villages that the 7200 PIH-supported CHWs serve is a 1000-fold increase; scaling from the three PIH-supported district to the 30 total in Rwanda is another 10-fold increase. This would put roughly two million rural homes (National Institute of Statistics of Rwanda 2012, p. 9) on the map.

5.6 *Potential Impact and Future Directions*

PIH successfully used GIS data to improve healthcare infrastructure in underserved areas (see section 1.3). Their team has integrated the house-level data from our pilot into their maps (see section 5.3). Such maps have the potential to improve PIH's operations further. The geography of catchment areas (see geospatial analysis in section 4.4) could be used to optimize the coverage of households, or as one of the criteria when nominating new CHWs. The higher resolution village maps could help the MOH plan and develop infrastructures, such as new drinking water wells and optimally placed small health outposts.

Furthermore, the location and density of the population could be the basis for future risk analyses that include a geographical component, such as flooding and mosquito-borne diseases. Linking health outcomes to geographical maps could inform future health interventions. For example, cases of malaria linked to population maps could inform the distribution of mosquito nets, or vaccination rates of newborns linked to the location of the closest health facility could inform the development of the primary health care system.

In epidemiology, social characteristics of the population are important. Diseases that spread through human contacts, such as Ebola, could be traced using relationship data (see social network analysis in section 4.4). Interventions could be planned by predicting how such viruses spread in a human network.

The same techniques could also predict how information spreads in a human network. Information inputted *into* the network influence behavioral change in communities (see section 2.4). Health interventions could be targeted at opinion leaders in villages that are socially central (as defined in section 4.4). Such central community members might spread interventions effectively throughout villages and as a result, improve health for the entire population. Social central-

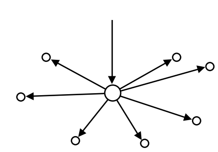

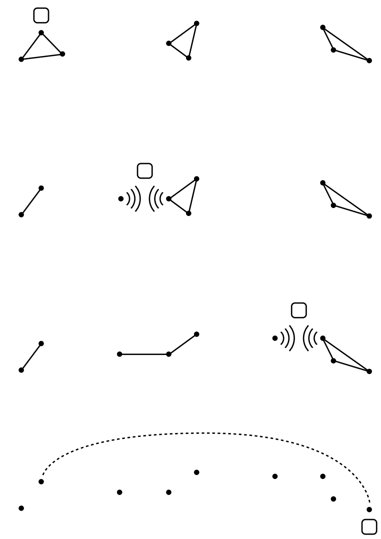

ity could be used in addition to spatial centrality (section 4.4 shows spatial and social centralities combined) as one of the criteria when nominating new CHWs.

With the established baseline map and network, additional functionalities are easier to add to the system. A dashboard component would be able to present collected data in real time. Two-way functionality to push survey questions to the mobile app could enable analysts to collect demographic or health-related data on-demand. The app could record walking movement to provide CHWs with statistics on their activities, and automatically map paths and roads. Furthermore, the app could offer an interface for the community to self-map social links to improve the social map.

A peer-to-peer communication function could enable community members to exchange information through the network. The system could enable zero-cost data transmission through Bluetooth and Wi-Fi from device to device. CHWs could distribute and route messages[15] or content by foot (see section 4.3 for observed instances of photo sharing and video distribution during our pilot) for the entire community through a store-and-forward mesh network.[16] Information could flow from person to person along the physical paths. A rich communication network in a region with little data connectivity could empower communities to create new services and increase local visibility.

[15] Video messages could easily be consumed and created by community members with any literacy-level.

[16] Small data packets could notify over the cellular network of big packets which are then delivered physically.

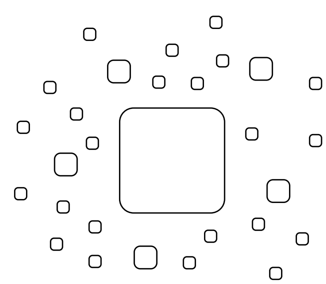

Applications situated at this intersection of the socio-spatial map (as defined in section 5.1) have the potential to go beyond rural communities. Human–machine collaboration offer a scalable method to create socio-spatial maps. New methods could be developed to derive social maps from spatial data, and spatial maps from social data. The high density of human activity in cities generates data for specialized maps at a rapid pace. Existing or new systems that use a socio-spatial map could be further studied to understand the dynamics between its elements – people, places, paths, and relationships – and their future applications better.

6 Conclusions

This thesis presents a collaborative human–machine crowdmapping approach to creating socio-spatial maps that represent both spatial and social aspects of communities. In particular, our system is designed to empower communities in unmapped regions of the world to make themselves digitally visible.

Our systems design combines (1) automated maps from satellite image analytics, (2) an intuitive mobile mapping app for novice users, and (3) social community survey data. We implemented this design and successfully tested it with communities in rural Rwanda. Our machine learning pipeline extracted building footprints from custom acquired high-resolution satellite imagery of that region. We show how local health workers mastered the use of the app with minimal training, integrated it into their daily routine, and mapped their entire community. The health workers enriched the machine-generated house-level map with GPS locations of households. The result is a detailed map with a structured mapping layer that is ground-validated. As far as we know, this approach of the two-way integration of top-down satellite data and bottom-up field data collected by local communities is novel.

The integrated approach of custom-built technology and human engagement resulted in a highly performing system: without prior smartphone experience, the volunteers rapidly mapped 100 percent of homes in the study's region at high accuracy. The app's ease of use and ability for users to develop new spatial insights through its design motivated them to map. One user commented "I knew where my households are, but the map changed the representation in my mind." (CHW 1.1, *Field Interview* 2017) Our analyses of user interviews and the collected data strongly indicate that with an appropriately designed solution, communities at the base of the pyramid can create state of the art maps and see potential in their continued use.

The design characteristics of our implemented social machine and field research provide practical insights that can be applied when designing new systems. For example, we show that user interactions can be simplified by focusing on a single action and delegating the rest to machine assistance. Furthermore, the thesis provides a case study of how using familiar concepts in interface design can make software accessible, and that aligning its use with existing behavior can lead to vibrant, self-sustaining systems. For the results to be actionable, it is essential to partner with a local organization. We shared the mapping data with Partners in Health who updated the map they use to optimize their health operations in Burera, Rwanda. We show how the elements of our system could scale beyond this region; in Rwanda alone, this could put approximately two million households living in rural regions on the map.

Our spatial and social analyses of the collected data offer new maps for this region. In combining the spatial mapping data with social survey data, we model a socio-spatial community network. Our analysis yields estimates for social centrality that identify influential opinion leaders, and the map representation enables observations about their spatial location in the village. Other forms of fine-resolution social relations data could be added into our model. The results have the potential to direct interventions in fields such as public health. To generalize this notion of social and spatial map elements, we organize them as mathematical graphs: a social graph consists of people connected by relationships, and a spatial graph consists of places connected by paths. We define a *socio-spatial map* as a map that combines *people, places, paths,* and *relationships* (fig. 5.1). In conclusion, this thesis presents a new approach for communities to become digitally visible, and provides a framework to organize socio-spatial maps.

References

"AAAS News and Notes" (2006). In: *Science* 312.5782, pp. 1891–1892. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.312.5782.1891. URL: http://www.sciencemag.org/cgi/doi/10.1126/science.312.5782.1891 (visited on 04/08/2017).

Apple (2015). *Model-View-Controller*. Cocoa Core Competencies. URL: https://developer.apple.com/library/content/documentation/General/Conceptual/DevPedia-CocoaCore/MVC.html (visited on 05/12/2017).

Azure Machine Learning Team (2015). *Microsoft Azure Machine Learning: Algorithm Cheat Sheet*. URL: http://aka.ms/MLCheatSheet.

Banerjee, A. et al. (2013). "The Diffusion of Microfinance". In: *Science* 341.6144, pp. 1236498–1236498. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1236498. URL: http://www.sciencemag.org/cgi/doi/10.1126/science.1236498 (visited on 05/20/2016).

Brandes, Ulrik (2001). "A faster algorithm for betweenness centrality". In: *The Journal of Mathematical Sociology* 25.2, pp. 163–177. ISSN: 0022-250X, 1545-5874. DOI: 10.1080/0022250X.2001.9990249. URL: http://www.tandfonline.com/doi/abs/10.1080/0022250X.2001.9990249.

Breiman, Leo (2001). "Random Forests". In: *Machine Learning* 45.1, pp. 5–32. ISSN: 1573-0565. DOI: 10.1023/A:1010933404324. URL: http://dx.doi.org/10.1023/A:1010933404324.

Brin, S. and L. Page (1998). "The Anatomy of a Large-Scale Hypertextual Web Search Engine". In: *Seventh International World-Wide Web Conference (WWW 1998)*. URL: http://ilpubs.stanford.edu:8090/361/.

Cancedda, Corrado et al. (2014). "Enhancing Formal Educational and In-Service Training Programs in Rural Rwanda: A Partnership Among the Public Sector, a Nongovernmental Organization, and Academia". In: *Academic Medicine* 89.8, pp. 1117–1124. ISSN: 1040-2446. DOI: 10.1097/ACM.0000000000000376. URL: http://content.wkhealth.com/linkback/openurl?sid=WKPTLP:landingpage&an=00001888-201408000-00014 (visited on 05/10/2017).

Chipchase, Jan (2005). *Understanding Non-Literacy as a Barrier to Mobile Phone Communication*. Tokyo: NRC (Nokia Research Center) Blue Sky. URL: http://research.nokia.com/bluesky/non-literacy-001-2005/index.html.

Chou, Yu-Kai (2016). *Actionable gamification: beyond points, badges, and leaderboards*. OCLC: 952061902. Fremont, CA: Octalysis Media. 499 pp. ISBN: 978-1-5117-4404-1.

Corbet, Jon (2009). *Good practices in participatory mapping*. International Fund for Agricultural Development (IFAD). URL: http://www.eldis.org/go/home&id=44098&type=Document.

Dennett, Daniel C. and Deb Roy (2015). "Our Transparent Future". In: *Scientific American* 312.3, pp. 64–69. ISSN: 0036-8733. DOI: 10.1038/scientificamerican0315-64. URL: http://www.nature.com/doifinder/10.1038/scientificamerican0315-64 (visited on 10/17/2015).

DigitalGlobe (2014). *WorldView-3 Data Sheet*. URL: https://www.spaceimagingme.com/downloads/sensors/datasheets/DG_WorldView3_DS_2014.pdf.

Dimagi (2014). *World Vision: Deploying MOTECH Suite to Support Global MNCH & Nutrition Programs*. Case Study, p. 9.

Dirk Robinson (2014). "The Skybox Vision". Proceedings of JACIE 2014 (Joint Agency Commercial Imagery Evaluation) Workshop. Louisville, Kentucky. URL: https://calval.cr.usgs.gov/wordpress/wp-content/uploads/14.018_Robinson-JACIE2014_public.pdf.

Garg, Vasundhara (2015). *Mapping Granular Household-Level Data for a Small Village in Odisha*. SocialCops. URL: https://blog.socialcops.com/inside-sc/announcements/mapping-granular-household-level-data-for-a-small-village-in-orissa/ (visited on 05/11/2017).

Google (2017). *Google Map Maker has closed*. Map Maker Help. URL: https://support.google.com/mapmaker/answer/7195127 (visited on 05/11/2017).

Google Earth Engine Team (2015). *Google Earth Engine: A planetary-scale geo-spatial analysis platform*. URL: https://earthengine.google.com.

Gros, Andreas and Tobias Tiecke (2016). *Connecting the world with better maps*. Facebook Code. URL: https://code.facebook.com/posts/1676452492623525/connecting-the-world-with-better-maps/ (visited on 11/04/2016).

GSMA (2016). *The Mobile Economy Africa 2016*. URL: http://www.gsma.com/mobileeconomy/africa/.

Ha, Linne (2015). *Crowdsourcing a Text-to-Speech voice for low resource languages*. Research Blog. URL: https://research.googleblog.

com/2015/09/crowdsourcing-text-to-speech-voice-for.html (visited on 05/11/2017).

Hancher, Matt (2013). *Only clear skies on Google Maps and Earth*. Google Lat Long. URL: https://maps.googleblog.com/2013/06/only-clear-skies-on-google-maps-and.html (visited on 05/12/2017).

Harvard Humanitarian Initiative (2011). *Satellite Sentinel Project: Evidence of Burial of Human Remains in Kadugli, South Kordofan*. Special Report. Harvard Humanitarian Initiative, p. 11. URL: http://satsentinel.org/sites/default/files/SSP%2018%20-%20Special%20Report%20Evidence%20of%20Human%20Remains%20in%20Kadugli%20082411.pdf.

Hendler, Jim and Tim Berners-Lee (2010). "From the Semantic Web to social machines: A research challenge for AI on the World Wide Web". In: *Artificial Intelligence* 174.2, pp. 156–161. ISSN: 00043702. DOI: 10.1016/j.artint.2009.11.010. URL: http://linkinghub.elsevier.com/retrieve/pii/S0004370209001404 (visited on 07/03/2016).

IDEO, ed. (2015). *The field guide to human-centered design: design kit*. 1st. ed. San Francisco, Calif: IDEO. 189 pp. ISBN: 978-0-9914063-1-9. URL: http://www.designkit.org/resources/1.

Kim, David A et al. (2015). "Social network targeting to maximise population behaviour change: a cluster randomised controlled trial". In: *The Lancet* 386.9989, pp. 145–153. ISSN: 01406736. DOI: 10.1016/S0140-6736(15)60095-2. URL: http://linkinghub.elsevier.com/retrieve/pii/S0140673615600952 (visited on 05/12/2017).

Knapp, Jake, John Zeratsky, and Braden Kowitz (2016). *Sprint: how to solve big problems and test new ideas in just five days*. First Simon & Schuster hardcover edition. OCLC: ocn913304205. New York: Simon & Schuster. 274 pp. ISBN: 978-1-5011-2174-6.

Knoche, Hendrik and Jeffrey Huang (2012). "Text is not the Enemy - How illiterate people use their mobile phones". In: *NUI workshop at CHI'12*. Austin, Texas, USA.

Kurtzman, Laura (2014). *UCSF*. Google Earth Engine Making Maps to Predict Malaria. URL: https://www.ucsf.edu/news/2014/09/116906/ucsf-google-earth-engine-making-maps-predict-malaria (visited on 11/04/2016).

LAND INFO (2017). *Buying Satellite Imagery: GeoEye, WorldView 1, 2, 3, QuickBird, IKONOS, Pléiades*. URL: http://www.landinfo.com/satellite-imagery-pricing.html (visited on 05/12/2017).

Lynch, Kevin (1960). *The Image of the City*. Cambridge, Mass. : M.I.T. Press, c1960. ISBN: 0-262-12004-6. JSTOR: HaydenLibrary-Stacks{NA}9108.L9871960. URL: https://search.ebscohost.com/login.

aspx?direct=true&db=cat00916a&AN=mit.001537571&site=eds-live&scope=site.

Mann, Julian (2014). *Introducing Skybox for Good*. Google for Nonprofits Blog. URL: https://googlefornonprofits.blogspot.com/2014/10/introducing-skybox-for-good.html (visited on 05/12/2017).

Marshall, Will (2017). *Planet to Acquire Terra Bella from Google, Sign Multi-Year Data Contract*. URL: https://www.planet.com/pulse/planet-to-acquire-terra-bella-from-google/ (visited on 05/12/2017).

Medhi Thies, Indrani (2015). "User Interface Design for Low-literate and Novice Users: Past, Present and Future". In: *Foundations and Trends® in Human–Computer Interaction* 8.1, pp. 1–72. ISSN: 1551-3955, 1551-3963. DOI: 10.1561/1100000047. URL: http://www.nowpublishers.com/article/Details/HCI-047 (visited on 04/12/2017).

Medhi, Indrani, Aman Sagar, and Kentaro Toyama (2006). "Text-Free User Interfaces for Illiterate and Semi-Literate Users". In: IEEE, pp. 72–82. ISBN: 978-1-4244-0484-1. DOI: 10.1109/ICTD.2006.301841. URL: http://ieeexplore.ieee.org/document/4085517/ (visited on 04/12/2017).

Ministry of Local Government (2006). *Ministerial Order Determining the Boundaries of Villages*. URL: http://rwandabar.org.rw/wp-content/uploads/2014/07/Ministerial-Order-determining-boundaries-of-Villages.pdf.

– (2010). *Bye Bye Nyakatsi*. URL: http://www.minaloc.gov.rw/index.php?id=514 (visited on 05/12/2017).

– (2011). *Administrative Units*. URL: http://www.minaloc.gov.rw/index.php?id=450 (visited on 05/12/2017).

Morrow, Nathan et al. (2011). *Independent Evaluation of the Ushahidi Haiti Project*. Evaluation. DISI – Development Information Systems International Ushahidi Haiti Project.

Munyaneza, Fabien (2016). "Using Community Health Workers (CHW) in low resource settings to map village level data". Technical Meeting, Improving Immunization Coverage and Reducing Inequities: Use of GIS in Immunization Programs. United Nations Children's Fund, New York City.

Munyaneza, Fabien et al. (2014). "Leveraging community health worker system to map a mountainous rural district in low resource setting: a low-cost approach to expand use of geographic information systems for public health". In: *International Journal of Health Geographics* 13.1, p. 49. ISSN: 1476-072X. DOI: 10.1186/1476-072X-13-49. URL: http://ij-healthgeographics.biomedcentral.com/articles/10.1186/1476-072X-13-49 (visited on 07/08/2016).

Murgia, Madhumita (2016). *FT's Seasonal Appeal: the volunteer mapmakers saving lives*. Financial Times. URL: https://www.ft.com/content/7a6df79c-c139-11e6-81c2-f57d90f6741a%60 (visited on 01/12/2017).

NASA (2017). *International Satellite Cloud Climatology Project: Cloud Climatology*. URL: https://isccp.giss.nasa.gov/role.html#DISTRIBANDCHAR (visited on 05/12/2017).

National Institute of Statistics of Rwanda (2012). *Rwanda Fourth Population and Housing Census*. Thematic Report.

Ntabathia, Jude (2015). "Mapjack: A Mobile based Wiki for Collaborative Map Making". Master of Science in Media Arts and Sciences Thesis. Massachusetts Institute of Technology. 70 pp. URL: https://dspace.mit.edu/handle/1721.1/98623.

Ntiamoah, Eva (2016). *Ghana Enters into a contractual agreement with ESOKO - Government of Ghana*. URL: http://www.ghana.gov.gh/index.php/media-center/features/2765-gender-ministry-partners-esoko-ghana-to-deliver-social-protection-for-the-poor-vulnerable-under-leap-programme (visited on 05/11/2017).

Nuwer, Rachel (2014). *The last unmapped places on Earth*. BBC Future. URL: http://www.bbc.com/future/story/20141127-the-last-unmapped-places (visited on 05/10/2017).

Pelechrinis, Konstantinos and Prashant Krishnamurthy (2016). "Socio-spatial affiliation networks". In: *Computer Communications* 73, pp. 251–262. ISSN: 01403664. DOI: 10.1016/j.comcom.2015.06.002. URL: http://linkinghub.elsevier.com/retrieve/pii/S0140366415002169 (visited on 05/11/2017).

Price, Rhiannan (2015). *Eradicating polio with the help of remote sensing*. DigitalGlobe Blog. URL: http://blog.digitalglobe.com/industry/eradicating-polio-with-the-help-of-remote-sensing/ (visited on 05/10/2017).

– (2017). *Eliminating malaria with the power of the crowd*. DigitalGlobe Blog. URL: http://blog.digitalglobe.com/crowd/eliminating-malaria-with-the-power-of-the-crowd/ (visited on 05/10/2017).

Saveski, Martin et al. (2016). "Human Atlas: A Tool for Mapping Social Networks". In: ACM Press, pp. 247–250. ISBN: 978-1-4503-4144-8. DOI: 10.1145/2872518.2890552. URL: http://dl.acm.org/citation.cfm?doid=2872518.2890552 (visited on 05/12/2017).

Schaad, Raphael (2017). *Visible Communities Community Health Worker Field Interviews*. In collab. with CHWs. Trans. by Fabien Munyaneza. In villages: (1) Mariko, (2) Kiraro, (3) Rusebeya, (4) Gisirwe. COUHES No. 1611763169, Self-mapping of Rural Villages in Rwanda by Community Health Workers Using a Smartphone App.

Schingler, Robbie (2017). *Planet Launches Satellite Constellation to Image the Whole Planet Daily*. URL: https://www.planet.com/pulse/planet-launches-satellite-constellation-to-image-the-whole-planet-daily/ (visited on 05/12/2017).

Singh, Prabhjot and Jeffrey D Sachs (2013). "1 million community health workers in sub-Saharan Africa by 2015". In: *The Lancet* 382.9889, pp. 363–365. ISSN: 01406736. DOI: 10.1016/S0140-6736(12)62002-9. URL: http://linkinghub.elsevier.com/retrieve/pii/S0140673612620029 (visited on 05/11/2017).

Snow, J. (1855). *On the Mode of Communication of Cholera*. John Churchill. URL: https://books.google.com/books?id=-N0_AAAAcAAJ.

Terra Bella (2013). *SkySat-1 First Light*. Terra Bella Blog. URL: https://terrabellatech.blogspot.com/2013/12/skysat-1-first-light.html (visited on 05/12/2017).

The WNDW Authors (2007). *Wireless Networking in the Developing World: A practical guide to planning and building low-cost telecommunications infrastructure*. Second Edition. Hacker Friendly LLC. ISBN: 978-0-9778093-6-3. URL: http://wndw.net.

Uddbäck, Sandra (2016). *Introducing the Mapillary Humanitarian Mapping Kit, in Partnership with Humanitarian OpenStreetMap Team*. Crowdsourced Street Photos. Mapillary is a service for crowdsourcing street level photos using smartphones and computer vision. URL: https://www.hotosm.org/updates/2016-07-27_introducing_the_mapillary_humanitarian_mapping_kit_in_partnership_with_hot (visited on 05/11/2017).

UNESCO Institute for Statistics (2016). *UNESCO Global Literacy Fact Sheet*. 38, p. 10. URL: http://uis.unesco.org/sites/default/files/documents/fs38-50th-anniversary-of-international-literacy-day-literacy-rates-are-on-the-rise-but-millions-remain-illiterate-2016-en.pdf (visited on 12/04/2017).

U.S. Geological Survey (2016). *What are the band designations for the Landsat satellites?* URL: https://landsat.usgs.gov/what-are-band-designations-landsat-satellites (visited on 05/12/2017).

Varshney, Kush R. et al. (2015). "Targeting Villages for Rural Development Using Satellite Image Analysis". In: *Big Data* 3.1, pp. 41–53. ISSN: 2167-6461, 2167-647X. DOI: 10.1089/big.2014.0061. URL: http://online.liebertpub.com/doi/10.1089/big.2014.0061 (visited on 04/22/2016).

Walter, Aarron (2011). *Designing for Emotion*. Brief books for people who make websites 5. OCLC: 760163428. New York, N.Y: A Book Apart. 104 pp. ISBN: 978-1-937557-00-3.

Wiedenbeck, Susan (1999). "The use of icons and labels in an end user application program: An empirical study of learning and retention". In: *Behaviour & Information Technology* 18.2, pp. 68–82.

issn: 0144-929X, 1362-3001. doi: 10.1080/014492999119129. url: http://www.tandfonline.com/doi/abs/10.1080/014492999119129 (visited on 04/12/2017).

World Bank (2017). *Atlas of Sustainable Development Goals 2017: From World Development Indicators*. DOI: 10.1596/978-1-4648-1080-0. The World Bank. 115 pp. isbn: 978-1-4648-1080-0. url: http://elibrary.worldbank.org/doi/book/10.1596/978-1-4648-1080-0 (visited on 05/10/2017).

World Health Organization, ed. (1989). *Strengthening the performance of community health workers in primary health care: report of a WHO study group*. Technical report series / World Health Organization 780. Geneva : Albany, NY: World Health Organization ; WHO Publications Center USA [distributor]. 46 pp. isbn: 978-92-4-120780-5.

Yuan, Jiangye (2016). "Automatic Building Extraction in Aerial Scenes Using Convolutional Networks". In: *arXiv:1602.06564 [cs]*. arXiv: 1602.06564. url: http://arxiv.org/abs/1602.06564 (visited on 04/22/2016).

Zegras, Christopher (2013). *Mobile Phones for Sustainable Mobility in the Global South?* url: http://czegras.blogspot.com/2013/04/normal-0-false-false-false-en-us-x-none.html (visited on 05/11/2017).

Credits for Photos

Raphael Schaad	Figure 3.9 (top left, top right, bottom left, bottom right), Figure 5.2 (top left, top right, bottom right), and all in Appendix E
Zach Both	Figure 3.1 (top left, top right, bottom)
Fabien Munyaneza	Figure 5.2 (bottom left)

All figures by Raphael Schaad unless indicated otherwise.

Terra Bella SkySat Public Ortho Imagery, RGB:
https://explorer.earthengine.google.com/#detail/SKYSAT%2FGEN-A%2FPUBLIC%2FORTHO%2FRGB
Terra Bella SkySat Public Ortho Imagery, Multispectral:
https://explorer.earthengine.google.com/#detail/SKYSAT%2FGEN-A%2FPUBLIC%2FORTHO%2FMULTISPECTRAL

Appendices

A Study Consent Form and Questionnaire

Interview and Study Consent Form
Self-mapping of rural villages in Rwanda by community health workers using a smartphone app

Before initial interview:
My name is Raphael Schaad. I am a graduate student at M.I.T., a university in the U.S. I have come here to test a new smartphone application for community health workers. I will be using the information from the interview and study to make it easier to design new smartphone applications for people in Rwanda and assess the feasibility of self-mapping of rural villages as part of my Master's thesis.

You were selected as possible participant because you are a community health worker in Kaganda, the region of my study. You should ask questions about anything you do not understand, before deciding whether or not to participate. Details of the interview and study:

- The interview will last about 30 minutes. I first will ask you some basic questions, and I then will ask you to draw the village you live in. If you do not want to answer any question or draw, you do not have to, and you can stop the interview at any time. I will then give you a smartphone and show you how to use the application to check off house visits in your catchment area. When you use the application, it can record where you are walking. You can disable this at any time. I will record our conversation to take notes. I will not share the recordings and delete them at the end of the project.
- I will ask if I can visit you at your home for about 1 hour in a few days. During the visit, I would watch you use the application, write notes about my observations, and take photo and video. You do not have to accept the request, and you can ask me to leave at any time.
- After 10 days I will visit you at your home for about 30 minutes to collect the smartphone and to ask a few final questions. The smartphone has to be returned after the study.
- With the smartphone, you will receive 500 RWD ($0.6) for each day of the study to charge the device to a total of 5,000 ($6). At the end of the study, you will receive 5,000 RWD ($6) to reimburse you for the time spent answering our questions. You will not receive any money for using the application. Your participation is completely voluntary. If you decide to stop the study early, the reimbursement will be prorated. What you say or do will not have consequences of any kind.

You will receive a copy of this form translated into Kinyarwanda. Any information that can be identified with you will remain confidential and will be disclosed only with your permission. Is there anything you do not understand? Would you like to participate in the study and interview? Thank you.

Name of subject: ________________________ **Signature:** ______________ **Date:** ________

After initial interview:
Would it be acceptable to share the results of the study and interview including photo and video of you in publicly accessible publications and presentations? You have the right to review any photo and video and have anything deleted for any reason. *[ask about each category of information as listed at the bottom of this page]*

The project will be completed by June 1, 2017. Any data not approved for publication will be stored securely and destroyed 1 year after completion of the study. Please contact me at schaad@mit.edu / +1-617-253-1000 at any time with any questions or concerns. If you feel that you have been treated unfairly, or if you have questions regarding your rights, please contact my supervisor at dkroy@media.mit.edu / +1-617-253-0596 or the Chairman of the Committee on the Use of Humans as Experimental Subjects (COUHES) at couhes@mit.edu / +1-617-253-6787.

For investigator use only:
Subject consents to sharing of ... *[check all that apply]*

☐ name	☐ user avatar picture	☐ recorded locations
☐ interview responses	☐ sketch maps	
☐ notes of visits	☐ photos of visits	☐ videos of visits

Initial Interview — Name of subject: ________________________________ **Date:** ________
[administered verbally with local translator]

Background information

Village name: _____________________________ Years lived here: ____

Gender *[observable]*: ________________________ Age: ____

Occupation: ______________________________ Family status: ___________________________________

Education: _______________________________ Languages: _____________________________________

Community health work (CHW)

1. For how many years have you been a CHW? ____ 2. How many households are in your catchment area? ____

3. Is there anything that makes your job difficult? ____

____

4. What would make your job easier? ____

____

5. How do you decide when to visit whom? ____

____

Technology

1. Do you have electricity? *Y / N*
- If Y, is it available every day? *Y / N* If Y, what do you use it for? _______________________________________

- If N, Where is the closest source of electricity? ____

2. What electronics devices do you have? *Battery radio, Solar cell, Cell phone*, *Smart phone*

Others: ____

3. Smartphone
- If Y, since when? ____________ What do you use it for? ____
- If N, have you used one before? *Y / N*

  - If Y, where? _________________________ Was it useful? _____________________________________

  - If N, what do you think it could be useful for? ____

Sketch map #1 — Name of subject: ________________________________ **Date:** ________
[handout sketchbook on blank page and take notes on the sequence in which the map is drawn]

I would like you to draw your entire village. Just describe the village to me, covering all the main features. You can start by marking where your house is.

____

____

____

____

____

Now, I would like to know what elements of the village you think are most important. They may be large or small, natural (lake, hill) or buildings (market, water well), but tell me those that for you are the easiest to identify and remember.

____

____

____

____

____

Please mark your catchment area and your peers' catchment areas. Please mark a few households.

____

____

____

____

____

For investigator use only:
Agreed date, time, and location of intermediate visit: ____

Agreed date, time, and location of final interview: ____

Final Interview — Name of subject: __________________________________ **Date:** ________
[administered verbally with local translator]

App usage

1. How many times per day did you use the smartphone? ____

2. How difficult or easy was it to use the application? *1–Very difficult / 2–Difficult / 3–Neutral / 4–Easy / 5–Very easy*

- If 1 or 2, what made it difficult? ____

- If 3 or 4, what made it easy? ____

3. What did you like the most about the application? ____

4. Did you show the application to others? *Y / N* If Y, what did they think? __________________________________

5. How was using the smartphone compared to a regular cell phone? _____________________________________

6. What else did you use the smartphone for? ____

7. When you started to see the map of your catchment area appearing, did you see anything surprising? *Y / N*

- If Y, what? ____

8. When you started to see the map of your catchment area appearing, did you learn anything new? *Y / N*

- If Y, what? ____

9. Once the map of the entire region is completed, what would you like to see or do with it? ___________________

____

Community health work (CHW)

What would make your job easier? ____

Social links

How often do two households interact with each other *(1–Never, 2–Annually, 3–Monthly, 4–Weekly, 5–Daily)*, if …

1. *1 / 2 / 3 / 4 / 5* … they have children attending the same primary school?
2. *1 / 2 / 3 / 4 / 5* … they are in the same traditional family ambulance group?
3. *1 / 2 / 3 / 4 / 5* … they are in the same economic cooperative (e.g. small fish or knitting sweaters)?
4. *1 / 2 / 3 / 4 / 5* … they go to the same market (that gathers people twice a week, e.g. Kinyababa or Rusumo)?
5. *1 / 2 / 3 / 4 / 5* … they go to the same shops (commercial center, e.g. Ruhinga or Rusumo)?
6. *1 / 2 / 3 / 4 / 5* … they go to the same place to get drinking water for their home?
7. *1 / 2 / 3 / 4 / 5* … they participate in the same cooking demonstrations?

Sketch map #2 — Name of subject: ________________________________ **Date:** ________
[handout sketchbook on blank page and take notes on the sequence in which the map is drawn]

I would like you to draw your entire village. Just describe the village to me, covering all the main features. You can start by marking where your house is.

____

____

____

____

Now, I would like to know what elements of the village you think are most important. They may be large or small, natural (lake, hill) or buildings (market, water well), but tell me those that for you are the easiest to identify and remember.

____

____

____

____

Please mark your catchment area and your peers' catchment areas. Please mark a few households.

____

____

____

____

How was drawing the second sketch map of the village compared to the first one? ________________________

____

____

Additional comments

Do you have any additional questions, comments, or suggestions for me? ______________________________

____

____

B Sketch Maps by CHWs

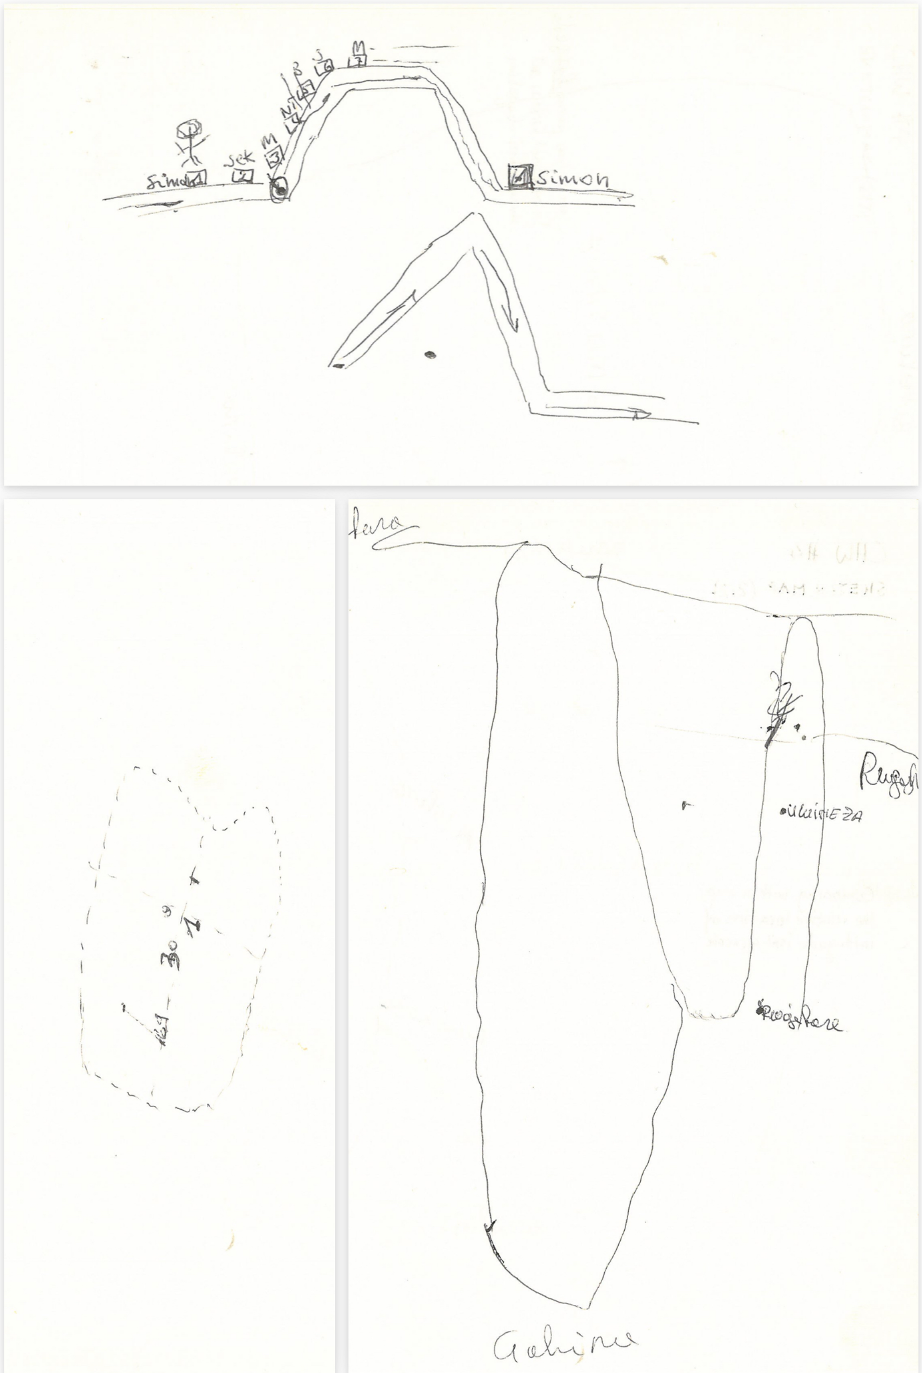

Examples of sketch maps by three different CHWs representing their catchment areas.

C Printed Map for CHW

One of twelve maps we created and printed during the field work for CHWs.

D Updated Map by PIH

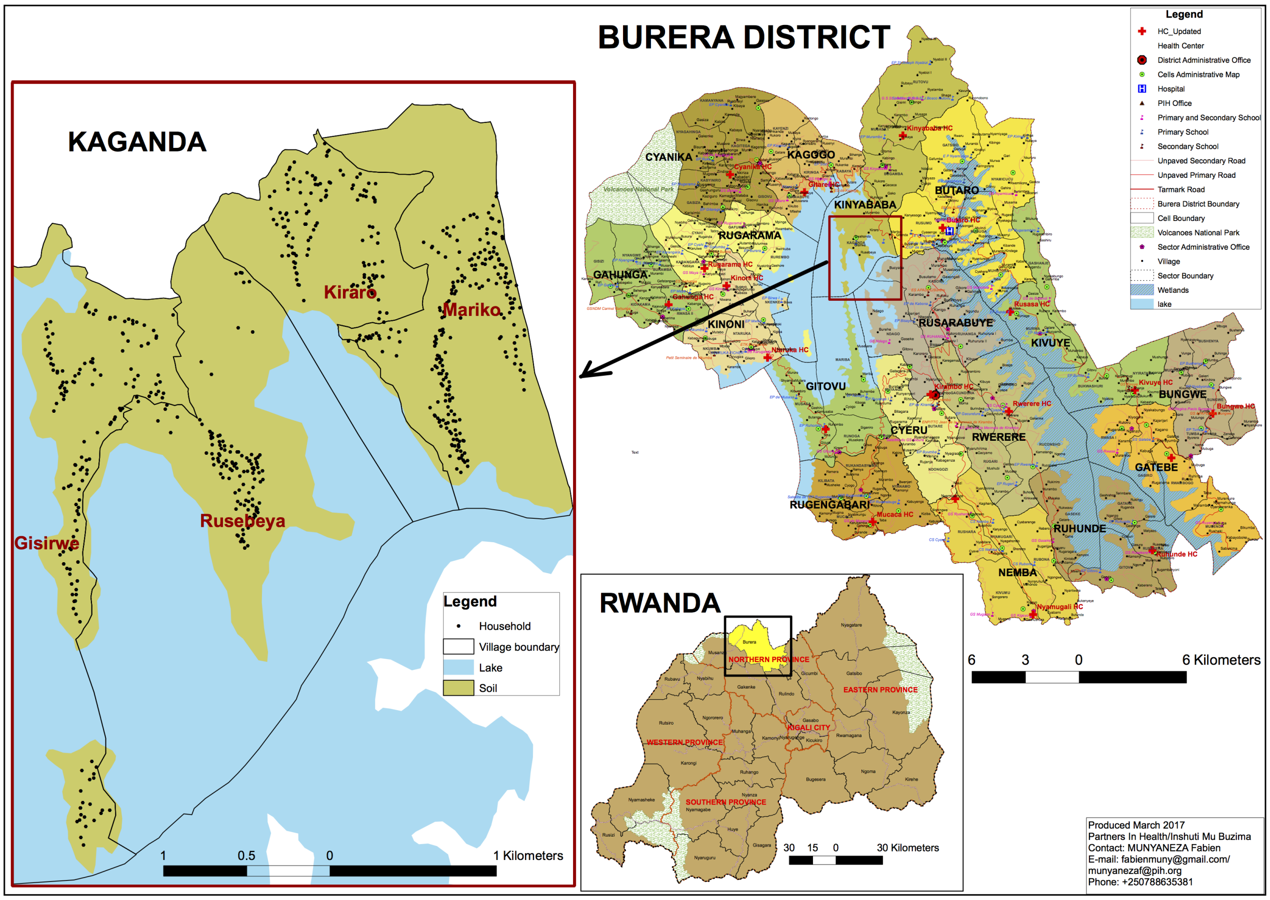

Burera district administrative map by PIH with household-level data from our pilot.

E Photos

Burera district administrative map by PIH on the wall of a public health center.

Rural villages in the mountainous region of Burera, Rwanda.

Locals in front of a typical single-family home with dried-mud walls and an iron sheet roof.

An electrified commercial center on a main road with a street side market.

A shop in a commercial center with mobile phone airtime and battery charging offerings.

A CHW sketching a map representation of their catchment area.

A CHW at their home, mapping the first household using the smartphone app.