Mobile Tools for Community Health Workers: An Application of Machine Vision for Mobile Phones

by

Xavier Andrew Soriano Díaz

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2018

Author .....
Department of Electrical Engineering and Computer Science
May 25, 2018

Certified by .....
Richard R. Fletcher
Research Scientist
Thesis Supervisor

Accepted by .....
Katrina LaCurts
Chair, Master of Engineering Thesis Committee

Mobile Tools for Community Health Workers: An Application of Machine Vision for Mobile Phones

by

Xavier Andrew Soriano Díaz

Submitted to the Department of Electrical Engineering and Computer Science
on May 25, 2018, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

According to a recent report from the World Health Organization (WHO) and the World Bank, as many as half of the world's population lack access to essential health services [1], a factor that contributes to the deaths of more than 5 million children every year. To tackle these problems, countries like India have set in motion a number of initiatives, such as that of Community Health Workers (CHW). However, CHWs feel disempowered as they often lack access to convenient and reliable tools to perform their measurements.

In this thesis, we first propose a number of mobile tools that make use of the smartphone camera, combined with machine vision and augmented reality (AR) to extract, collect and analyze data from the image of the child, and provide the CHW with an accurate, faster and automated means of performing Neonatal and Child Health (NCH) measurements, while at the same time, revealing relevant feedback about the health of the patient. The basic measurements include anthropometric information such as Length, Weight and Middle Upper Arm Circumference (MUAC).

For initial deployment and field testing of the technology, we partnered with the Public Health Foundation of India (PHFI), which implements much of the training for their accredited social health activists (ASHA). Since the ASHA health worker program is a government program, we obtained approval from the Government of India and conducted a pilot usability study with 13 ASHA workers in the New Delhi slums using our technology, between the months of May and September of 2017. For the study, we created Baby Naapp, an Android application that served as an interface for the ASHA workers to create basic profiles with socio-economic information of their patients, and to perform calls to each of our measurement tools. The findings of this study helped to refine our algorithms and the user experience of our apps.

Additionally, we have also explored the use of multi-spectral sensing to capture essential physiological measurements. We have developed an Android mobile application which uses a thermal camera module to automatically measure respiration rate (RR) and respiration rate variability (RRV). This app makes use of machine vision algorithms to detect a human face and then measure the temperature fluctuations in

the nostrils to approximate air flow rates and calculate timing parameters.

Future research will improve the accuracy of the anthropometric measurement tools by refining the machine vision algorithms and the hardware used to capture the readings, and will also include the thermal respiration app into the NCH kit. The data capturing algorithm will be used to develop a jaundice screening tool, and the data structures created to organize and analyze measurements in the NCH kit can be incorporated into the other verticals of our lab, such as the pulmonary and cardiovascular screening tools.

Thesis Supervisor: Richard R. Fletcher
Title: Research Scientist

Acknowledgments

I would like to start by thanking my advisor Rich Fletcher, who constantly guided my work and pushed me to explore more. Rich is the kind of person who understands that with great power comes great responsibility, and has decided to devote his skills to the service of those who need it most, and in the name of those who can't. He is the kind of person that will ask you to stay late at the lab to meet a deadline, but he will stay there with you working at your side. Rich is a great scientist and a great human, and I am very grateful for the opportunity that he gave me to join his group.

I want to thank the MIT Tata Center, for its generous support with sponsoring my work, and its constant guidance to contextualize our tools by really understanding the pain points of our users. I want to specially thank my Tata mentor Nevan Hanumara. While Rich kept pushing me to explore, Nevan kept asking me to focus, helping me to find more balanced and realistic expectations.

I want extend my gratitude to the Department of Electrical Engineering and Computer Science and the mentors that I have met from there, specially the people that have trusted in me. I want to thank Hal Abelson for giving me the chance of TAing his Mobile Development Class. Without such credential, I would not have qualified to join the Mobile Technology Group. And of course I want to thanks Chris Terman, who supervised my undergrad advance project. By believing in me and in my idea, Chris help me recover the self-confidence that I greatly needed.

I want to extend a great thanks to all the researchers that participated in the major studies of this work. Thanks to the researchers in the Public Health Foundation of India, especially Dr Suparna, who so diligently contributed with their efforts for a successful study. Thanks to the S-VYASA Yoga University researchers, who hosted me and have proven to be fantastic and always available research partner. A big thanks to Honey Bajaj, who worked by my side during my first year, designing the AR targets and preparing all the materials to successfully kick-start the ASHA study. And to finish this block of gratitude, I want to thank my group peers who made the exhausting days in India more fun and exciting.

I would like to once again thank my parents. With their love, teachings and support, they have shaped me into the person that I am now. I want to thank my older sister, Desiree, who convinced me to give it a shot and apply to MIT. I literally would not be here without her motivation and example. I want to thank my little sister, who was only a little kid when I left home, but has grown to become one of my most trusted advisors. I want thank Dome, my niece, whose memory always brightens my day. Even in her absence, the lovely memories of growing up together, our senseless but funny arguments and our crazy conversations, always bring a smile to my face. I know she is rooting for me from heaven, but I miss her dearly.

A gigantic thanks to my beloved Alkistis and her lovely family, who always keep me in their thoughts. Alkistis has not only been my loyal companion during this stage of my live, as a matter of fact I would not had been able to complete this thesis without her. Most recently, she helped me by proofreading this work, generating figures, and fixing formatting issues. She was also always ready to help with any design problem that arose during the work of my master.

I will end by thanking my God. I believe that all I am, was and will be is because of His grace. I offer this work and my sacrifice to him, my Lord Jesus Christ.

E Somnio Ad Astra

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Community Health Workers: Context and Latent Needs

[1]

In many developing countries, child health screening is often performed by community health workers (CHW) or front-line workers, which can be part of various public or private initiatives. These community health workers, nurses and midwives, are the first and often the only link between the health care system and billions of people that belong to the most marginalized populations.

CHWs are critical in settings where the overall primary health care system is weak or inaccessible, a problem that often requires them to go door-to-door, providing education, counseling, health screening, and referrals for institutional check-ups. In India specifically, the National Health Mission operates a community health worker program which consists of more than 700 thousand ASHA workers (Accredited Social Health Activist). These local women are nominated from their own community and receive some basic training for home-based newborn care (HBNC), and Reproductive, Maternal, Newborn, Child, and Adolescent Health (RMNCHA) services. The ASHA

[1]Some of the information in this work and in this particular chapter was originally published and presented at the IEEE GHTC 2017 as *Development of smart phone-based child health screening tools for community health workers*[2] on which the author of this thesis participated as a coauthor.

program was initially implemented in rural areas throughout India but is now being adopted in many poor urban slum areas as well [3].

Figure 1-1: ASHA worker on her way to a slum

Despite the existence of international guidelines, the quality of child health screening can vary widely, even across parts of the same country. Members of our team have spoken with doctors and health functionaries in developing countries, and some have even traveled to rural clinics in over 14 states in India, and observed many challenges, including the following:

- Lack of tools: Although a basic set of tools and medicines are supposed to be provided to the Indian ASHA workers [4], in practice many of these tools are missing or are of low quality.

- Lack of digital records: In most areas, data is collected manually and recorded on paper. This is not only prone to error, but also creates a more generalized pain point of redundant data that is entered over and over again, but never established as a shared resource.

- Limited expertise and education: While ASHAs are required to have a minimum education and be good with numbers, the condition may be relaxed if no other candidate is available. This limited education added to their limited clinical expertise results in a slow integration and adoption of any technology.

ASHAs also face socio-cultural prejudices and taboos, lack of a fixed income (since their remuneration is incentive/referral based), and heavy work loads that require them to cover great distances in a single day. Taken together, these challenges often result in a lack of capacity to detect early signs of disease or pathological conditions. More importantly, the lack of quantitative tools limit the ability for an uneducated health worker to make confident evidence-based medical referrals and motivate a child's family to take action or instill health-seeking behaviors. Generating a demand for health services is ultimately limited by the ability of a community to properly assess a child's health.

Figure 1-2: Rural Clinic in Punjab India

1.2 Existing Tools used by Community Health Workers

The National Health Mission of India lists the items in table 1.2 as the components of the HBNC kit for ASHA worker, as well a number of basic medicines that can be found in [?]

	List of Items in ASHA Equipment Kit
1	Digital Wrist Watch
2	Thermometer
3	Weighing Scale (for newborn babies)
4	Baby Blanket
5	Baby Feeding spoon
6	Kit Bag
7	Communication Kit
8	Mucous Extractor

Table 1.1: List of Items in ASHA Equipment Kit

Additionally, since as early as 2013, there have been talks of the government providing the ASHAs with mobile phones [5]. However, in practice the components of the ASHA kit are decided at the district level with some kits being more comprehensive than others. It is unclear how advance such deployment is, but as far as our experience goes, we are unaware of ASHAs with government issued smartphones. Nevertheless, the rapid reduction of production costs of smartphones, added to their exponentially increasing processing power and the ever increasing market penetration in development countries, make the prospect of ASHAs having their own smartphones soon far from an unreachable reality.

1.3 Emergence of mHealth Tools

A review of current and potential uses of smartphones in healthcare published by Stanford researchers on 2012 [6] identified five main categories: monitoring, health apps for untrained patients, communication, and reference apps for medical practitioners. Another systematic review by researchers in University of Missouri essentially

covered the same categories, but made emphasis on who the user for the application was. Both studies, had a selection criteria that focused on the functionality of smartphone software, sensors, and peripherals. The selection started from a pool of more than 2000 studies that got eventually reduced to around 60 papers, all ranging in between the years 2001 and 2012 and with a noticeable increase starting in 2010. The ubiquitous penetration of smartphones since then, make the 6 years that have passed seem like a few lifetimes in technology time. However, their analysis stays relevant and it can be expanded to the following categories.

In the category of patient monitoring, the smartphone sensors can be valuable resources for people who are unable to physically attend hospitals and clinics. Some examples include the following:

- Applications have been used to track the recovery of patients that have had a stroke using Bluetooth sensors that connect to the phone from the shoes.

- ECG sensors can be plugged in the smartphone via usb, or paired via Bluetooth and Wi-Fi to communicate real time heartbeat monitoring.

- GPS sensors can track the location of people with Alzheimer and dementia and determine if the patients have become confused, lost, or if they are in need of help.

- Systems like 99DOTS allow for medical adherence control using simple phone calls.

Other health apps are not necessarily used for monitoring of diseases. They are rather oriented to casual and seamless tracking of health indicator metrics in the daily routine of a person. These apps include calorie counters, step counters, sleep quality assessment, wellness and nutrition which gained traction in the west with the introduction of commercial products such as Fitbit and Apple Watch.

Additionally, smartphones can serve their primary function: a means of communication. Community health workers can consult physicians via text, image or video

messages embedded in custom apps that can be used for telescreening and telediagnosis [7].

The large storage available nowadays, combined with internet connectivity, allow medical practitioners to have immediate access to an enormous amount of reference information. Solutions like Epocrates, an athenahealth service, even claim that their app can save physicians 20 minutes of time each day.

Finally, small companies such as Dimagi and Mobilitas now offer software platforms that help simplify the task of survey-based data collection, support case management and have made fantastic advances in integrating with electronic medical record systems.

1.4 Shortcomings of existing mHealth Tools and Current Needs

Despite the existence of all these mobile platforms and their increased popularity in the developed world, the penetration of sensor-based data collection in the field is still minimal. For electronic medical record systems, all data collection normally requires manual data entry (e.g. typing in the glucose reading, or the baby's height and weight). Given this context, we have identified the need to create a next-generation of mobile tools that can make use of the intrinsic capability of smartphones to automatically make measurements without the need for a manual data entry.

Chapter 2

Machine Vision as a proposed solution

Machine Vision is a broad term that describes computer-based analysis of camera images, generally in real-time. In the past, Machine Vision algorithms have been used in a wide range of activities from improving quality of manufacturing and assembling, robotics, self-driving vehicles, to medical imaging. In the case of medical applications, processing power has traditionally been a limiting factor when developing low cost tools. However, with the advent of the smartphone and its explosive growth, Machine Vision is now being packed in these small gadgets, and becoming increasingly present even in daily communication apps.

A particular branch of Machine Vision is Augmented Reality (AR), which was recently popularized by the game Pokemon Go. In formal terms, AR is an overlay computer-generated sensory input such as graphics, video or sound over the real-time world environment. Development kits for AR mobile applications support several main types of AR, such as markerless based and maker based AR. For markerless based AR, a smartphone uses geolocation, accelerometer readings and other location estimators, combined with spatial recognition algorithms to understand the composition of a scene and place, and maintain AR objects in a particular position and orientation in the camera view. On the other hand, marker based AR uses special markers, known as targets, which are pre-established and known by the developer

and that are used to perform tracking and to learn the information of the physical characteristics of the scene.

Figure 2-1: Examples of AR — (left) Visor of pilot, with superimposed flight info; (right) Pokemon Go

Current smartphones contain many sensors and embedded technologies that can be applied to healthcare. Arguably, the most powerful of these technologies is its camera, with phone manufacturers deeply focusing on improving its capabilities over the past few years. Coupled with the also ever-improving processing power that allows to perform Machine Vision algorithms and AR, limitations from 20 years ago slowly vanish and the smartphone camera can be used as a powerful tool to automate and complement many tasks that are performed by health workers and doctors.

2.1 Prior Use of Machine Vision and AR in our Group for Biomedical applications

Just as in industrialized factories, Machine Vision has the potential to be used as a tool to augment or even automate health related procedures and with the continuous improvements of imaging hardware, it is increasingly penetrating the health research sector with clever and ambitious techniques and tools.

In our group, we have used Machine Vision to measure a patient's heart rate by analyzing small changes in his skin color. This technology is known as video plethysmography and it was partially developed at the MIT Media Lab and the Computer

Science and Artificial Intelligence (CSAIL). We have extended this technique to a smartphone application that automatically detects the person's face and samples the color pixels from a segment of the image [8]. However, testing results revealed that video plethysmography is not practical for a health worker using a hand-held phone as the technology is too sensitive to motion artifacts. Nevertheless, it has shown to be a promising technique for performing non-contact measurement of babies in situations where the phone can be mounted on a tripod or as part of a baby crib.

Additionally, our group also conducted a study that uses colorimetric tests for performing point-of-care health screening using blood samples [9]. The challenge addressed was that due to human color perception, this type of test results are difficult to quantify. For this purpose, a mobile app that performs the color comparison was introduced. The app uses AR to automatically identify the position and orientation of a printed test, which includes a target to facilitate the tracking algorithm, as well as a reference color palette to enable color calibration and compensation for environment settings. The solution is highly scalable and low cost, as it only requires tests printed on paper and a mobile app that can be easily installed on any Android phone. The study also suggested that measurements captured using mobile phone performed favorably when compared with manual measurements.

Figure 2-2: Examples of Prior Use of AR in our group — (left) AR App to automatically digitize the reading from a Peak Flow Meter; (right) a color test that uses AR to rectify the image and sample pixels from the correct colors

2.2 Machine Vision for Anthropometric and Physiological Measurements

In this thesis, we explore the use of Machine Vision to extract anthropometric and physiological measurements from babies and children, and other potential applications into global health. As with the printed diagnostics example, Machine Vision is an extremely convenient tool to automatically track regions of interest and liberate the lay users from the often challenging task of accurately aiming the camera to the desired location. In addition, AR is used to create live graphical overlays that provide critical feedback about the measurements in the scene.

Chapter 3

Anthropometric Measurements

3.1 Motivation

3.1.1 Importance of Anthropometric Measurements

The leading causes of death for children under 5 years of age are infectious diseases (e.g. pneumonia, diarrhea, malaria, measles) and malnutrition. Recent reports from UNICEF [10], state that nearly half of all deaths of children under the age of 5 can be related back to malnutrition, as it increases the frequency and severity of common infections, and usually delays recovery times. Even though the cases of children stunning have been slowly decreasing, malnutrition or wasting still affected the lives of 51 million children on 2017. While it was possible for many of these cases to see recovery, some of the damage may never be reverted as poor nutrition during the first 1000 days of life can lead to stunted growth and future impaired cognitive performance.

Given that newborn babies are generally considered a vulnerable population, anthropometric and physiological measurements begin at birth with neonatal screening, to identify potential congenital disorders or developmental challenges, such as prematurity. However, after the first few days of life, a child's health assessment generally shifts to monitoring the child's growth and nutrition status as well as to identifying signs of potential infectious diseases. The World Health Organization (WHO) has

published guidelines for child health assessment in an effort to standardize practice and improve data quality [11]. These guidelines are used by primary care clinics, community health workers and health camps worldwide.

While the assessment tools and metrics for child health vary as a function of socioeconomic status and type of clinical facility, certain fundamental metrics are common worldwide, even in low-resource settings. These traditional metrics include anthropometric information such as: height or length, weight, and middle-upper arm circumference (MUAC). Unfortunately for front-line workers, it is unfeasible for them to carry a bulky infantometer or an electric digital scale, which are considered the gold standard tools to collect anthropometric measurements. Instead, the length measurement is commonly not recorded, while the other metrics are collected with rudimentary tools and logged by hand in paper journals and notebooks.

3.2 An AR framework for anthropometric data collection

In an effort to overcome the lack of access to robust anthropometric measurement tools that CHW face everyday, we have developed a suite of Android Apps that make use of the smartphone camera combined with machine vision and augmented reality (AR) to extract, collect and analyze anthropometry from camera frames containing the scene with a child. We focused on the three measurements mentioned in the previous section: length, weight, and MUAC.

3.2.1 General Implementation Details

The mobile operating system of choice is Android due to its high penetration in smartphone market of the developing world. In addition, we are making use of open source libraries that already implement known Computer Vision algorithms and the basis of AR:

- OpenCV - A widely used and constantly maintained library of programming

functions mainly aimed at real-time computer and machine vision. The capabilities of OpenCV are extensive: from simple tasks as converting and image into grayscale, to providing a whole pipeline for training a custom object detector. While OpenCV algorithms require the freedom and speed that only languages like C and C++ can provide, libraries for Python and Java are massively supported. In particular, we have made use of various versions of the OpenCV software development kit (SDK) for Android.

- Vuforia - Formerly developed by Qualcomm, and now owned by PTC Inc, it is a SDK that enables the creation of Augmented Reality applications. It supports iOS, Unity and Android development and it is complemented by a web-based platform on which users can add their own custom image "targets". Features are extracted from the targets and a set of feature files is generated to be incorporated into the user's project, allowing for automatic recognition and tracking of the targets in real-time.

- OpenGL - A 2D and 3D graphic rendering API, which is supported by Android. Vuforia samples use OpenGL as the default renderer of AR objects.

We have augmented the already existing tools with simple stickers and patterns that serve as AR targets. Different AR targets designs and algorithms were developed to meet the specifications and constraints of each anthropometric measurement, but it is possible to abstract a common framework or pipeline that is followed by our three tools.

1. The phone automatically tracks the AR target and extracts information about its location in the camera frame

2. The target location is used to rectify and extract the intended region of interest.

3. A measurement-specific machine vision algorithm is executed to extract the required information.

4. Readings from every frame are aggregated to display real-time feedback about the measurement.

5. When the measurement has converged, the user is notified and can proceed to have the anthropometric reading captured and saved.

The design process, algorithm and validation of each individual tool will be explained in detail, but first, the following sections will go over the common elements of the pipeline.

3.2.2 Geometric Rectification

On Feature Extraction and Homographies

As briefly mentioned in a previous section, Vuforia generates feature files that are used to find and track targets in the camera frame. While the library documentation does not specify which feature extraction algorithm is used for proprietary reasons, there has been a number of successful features used for object detection, and object tracking. Scale Invariant Feature Transform (SIFT) is an example of such algorithm [12]. It presents an attractive way of identifying distinctive scale and rotations of invariant features from images, that can be used to match objects in different scenes, with automatic panorama stitching being one of its current most popular applications. For automatic panorama stitching, features are extracted from the images that share parts of the same scene. These images are obtained from several frames and different rotations of the camera. The nature of such rotations determine that the partial scenes can be related by a particular set of homographies [13]. A homography is a transform that relates two projection planes that share the same center of projection or are images that are viewing the same plane from a different angle. It can be represented as a 3×3 matrix, relating points of the two planes expressed in homogeneous coordinates. Homographies can be solved with at least four points, but in the case of panorama stitching, more than four potential matches will likely be found. For these cases, a combination of RANSAC (random sample consensus) [14] and least squares optimization is also shown in [13].

Target Tracking and Perspective Correction

Homographies are also, very likely, the basis used by Vuforia to understand the position and orientation of planar targets. Features extracted by its algorithms are used to estimate a projection matrix that OpenGL uses when rendering the AR information. From there, we are able to extract the locations of the four corners of our target. However, these matrix and corners exist in a frame of reference of OpenGL, which has its origin in the center of the camera frame and the ranges of its visible axes are normalized, i.e. they span values from -1 to 1. To be able to process the image in the frame, the corners and the projection matrix need to transformed into the frame of reference of OpenCV, which has its origin on the top left corner, with its axes spanning the pixel values of the screen resolution, the positive x axis pointing towards the right and the positive y axis pointing downwards. Because OpenCV image manipulation only allows for rectangular cropping, the outer-most coordinates of the region of interest (ROI) are used to generate a preliminary Bitmap that contains that ROI in its totality. It is important to take note of the OpenGL coordinates of this rectangle's top left corner $(\text{offset}_x, \text{offset}_y)$, to keep a reference of the OpenGL origin. The corner conversion is given by a set of equations of the form:

$$p_{openCV,x} = \text{int}\left(\frac{p_{openGL,x} + 1}{2} \times width_{pixels} - 1\right) - \text{offset}_x$$
$$p_{openCV,y} = \text{int}\left(\frac{1 - p_{openGL,y}}{2} \times height_{pixels} - 1\right) - \text{offset}_y$$

Subsequently, an inverse projection matrix needs to be calculated to obtain a rectified version of the region of interest from the camera frame. The OpenGL projection matrix is again inconvenient, as it contains an inferred value for the Z coordinated in the scene, i.e. an estimate of the depth component in a 3D world. A much simpler alternative is to use the OpenCV coordinates of the corners of the target, and find the perspective transformation that maps it to a strategically scaled version of its original image. Fortunately, such operation requires exactly 4 correspondent points, and given

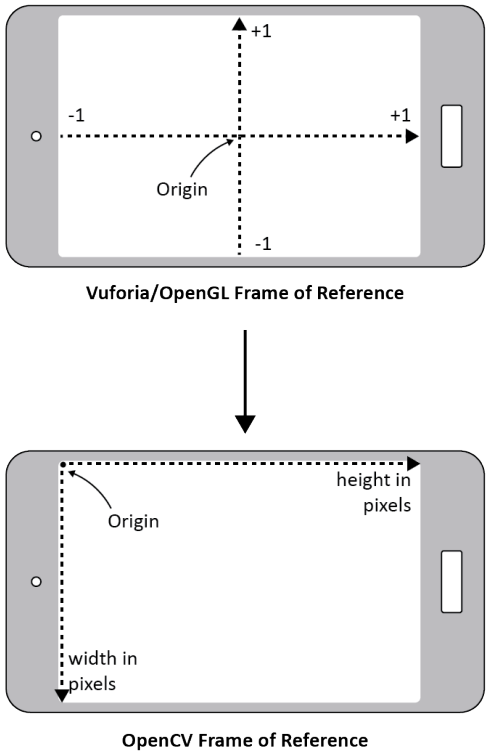

Figure 3-1: Schematic of OpenGL and OpenCV Frames of Reference

that the aspect ratio of the original target are known by the user [1], it is possible to easily obtain a rectified ROI for which its top right matches the origin of the OpenCV frame of reference. In other words, given the desired width w_t and height h_t of a rectified target and for the set of corners $(P_{top,right}, P_{top,left}, P_{bottom,left}, P_{bottom,bottom})$ in the OpenCV frame of reference, we need to find the perspective transform T that satisfies:

$$(P_{top,right}, P_{top,left}, P_{bottom,right}, P_{bottom,left}) \xrightarrow{T} ((0,0),(w_t,0),(w_t,h_t),(0,h_t))$$

OpenCV contains a Image Processing package named *ImgProc* which contains a convenient method *getPerspectiveTransform* that takes as inputs the source and destination points and generates our desired matrix T. *ImgProc* also counts with the

[1] In fact, the dimensions of the physical target are established by the user when creating it and are necessary to generate the feature files from the Vuforia Web Portal

method *warpPerspective*, which is used to finally obtain a rectified version of our region of interest. It is important to note that the pair of four points sets need to be correspondingly sorted when entered into *getPerspectiveTransform*.

3.2.3 Handling Continuous Data Feed

After each anthropometric tool has executed its custom algorithm to extract the measurement, a new measurement value is generated if there is a measurement present in the frame. With current frame rates it is possible to obtain between 5 and 10 measurements every second. Such continuous data feed requires careful handling to always maintain the most accurate estimate of the measurement possible.

A Sorted FIFO Queue of Measurements

For these purposes, a custom data type was developed and descriptively named *SortedMeasurementQueue*. While Queues are by definition FIFO (First In, First Out) data buffers, it is important to emphasize its correct behavior, as the *Sorted* label in its name may be a cause of confusion. The data type behaves like a normal *Queue*: as expected, older measurements are less important, so once the *SortedMeasurementQueue* reaches its full capacity, they are discarded in favor of newer ones. The special thing about this data type is that it also keeps track of the order of its elements to efficiently provide the median value. Because the maximum size of the *SortedMeasurementQueue* ought be kept smaller than a 2 second interval in order that preserve the validity of its contents, a simple *LinkedList* is sufficient to rapidly insert a new value in its correct location. If the queue is not full the value is coupled with the current size of the queue, while the "popping" pointer remains pointing at the first value that was entered. If the queue is at capacity, the new value is coupled with the index of the current position of the "popping" pointer, and the pointer increments by 1 modulo the size of the queue. In this way putting and popping values from our data type are $O(n)$, but when the queue becomes full and because of its relatively small size, one can think of the complexity of these operations as $O(1)$. The median value of

the *SortedMeasurementQueue* is saved as the measurement for the current frame, and it is kept in an extra queue that holds the values from about the last 2 seconds of measurements. This queue effectively holds a median filtered version of the signal of raw measurements.

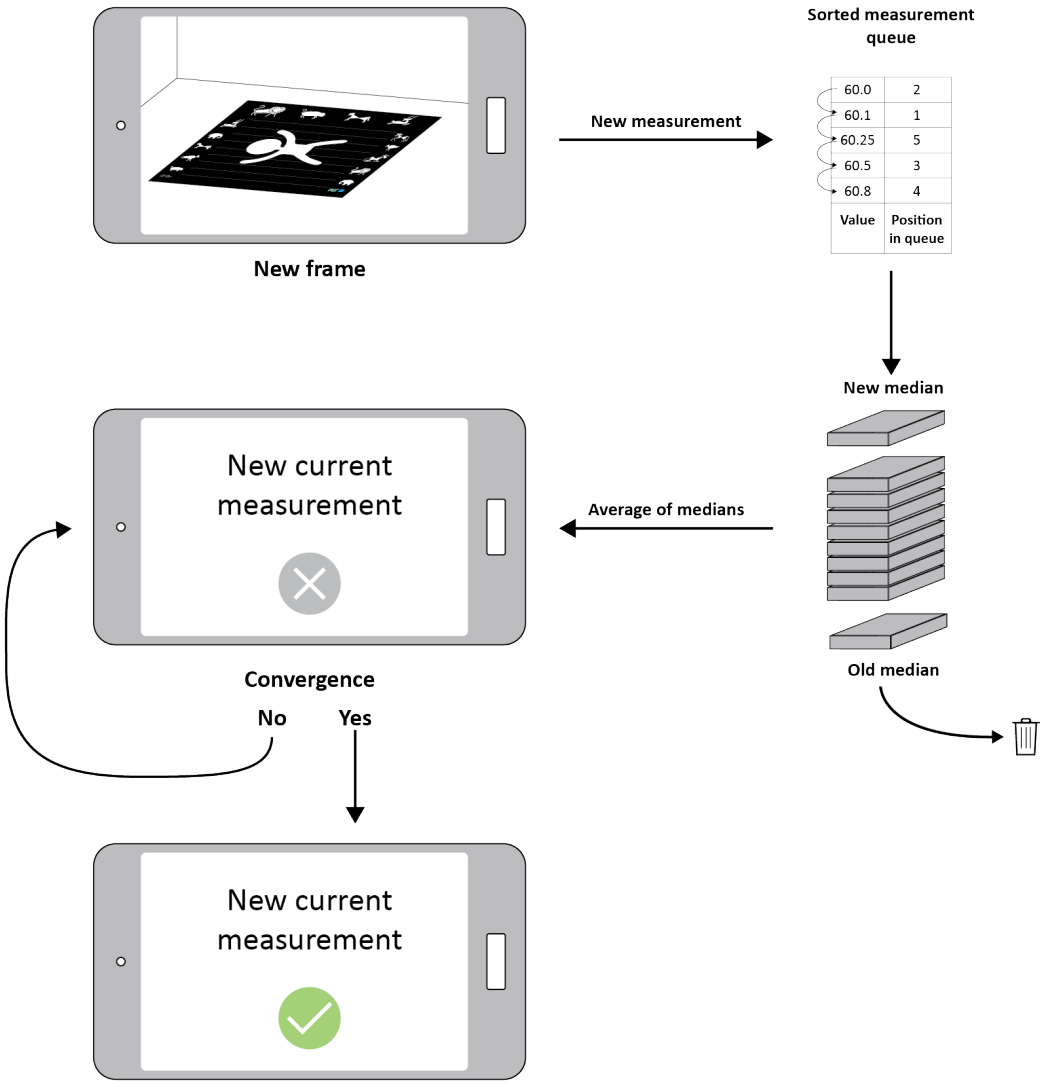

Figure 3-2: Continuous Data Feed Handling Workflow

A simple criterion for convergence

The measurement that is presented as the current one to the user is the average of the queue of medians. This value can be thought of as rolling average of the median-

filtered version of the raw signal. If the difference between the new median and the previous average is less than a user specified threshold, the new current average is labeled as stable. If the updated current values are stable for around two seconds, the algorithm determines that the measurement has converged, and allows the user to save the measurement. The whole pipeline is described in Figure 3-2.

3.3 Measurement of Length

3.3.1 Design of AR Target

Figure 3-3: Final Version of Baby Blanket used for the length measurement. The animal drawings constitute an AR target.

The tool used for the height measurement is a rectangular blanket that contains animal drawings on three of its sides and upon which the baby is placed in the direction of the horizontal guide line. The most recent design of the blanket can be seen in Figure 3-3, while older iterations are discussed in section 4.4.2. Besides creating a beautiful ornamentation for the tool, the animal drawings serve as a very feature

dense AR target that allows to automatically correct for geometric and parallax distortion, unlike other camera-based approaches that require significant calibration or sensor measurements [15, 16].

The choice of having the drawings on only three sides was not random. For this technique to work properly, the ASHA needs an aid to ensure that the legs of the baby are extended, and the toes making a 90 degree angle with respect to the ground. The mother is expected to provide such help, by holding the baby, and a target-less side of the blanket is left on purpose so that the presence of the mother does not interfere with the tracking of the baby.

The dimensions of this blanket are 98×78 cm and can measure infants of up to 68 cm, which is the average for 7 month-old babies.

3.3.2 Length Algorithm

The height extraction algorithm is simple and it is based on blob detection. Such simplicity, requires some heavy lifting performed at the preprocessing stage. The whole process develops as follows.

1. A ROI is selected inside the rectified image. The ROI is an inner rectangle of the blanket that does not contain the target. The boundaries for this rectangle are selected empirically.

2. The contrast of the image is increased to reduce shadows generated by the baby or the arms of the users.

3. The image is transformed into grayscale for blob detection based on intensities.

4. A Gaussian filter is applied to remove small noisy artifacts.

5. Morphological operations are applied. This step is normally as preprocessing in segmentations applications such as brain segmentation in MRIs, to remove small connections that probably caused by noise. For our case, we used them to remove noise indeed, but very importantly to remove the horizontal guide

lines. This can easily be achieved by an application back to back of erosion and dilation [17] with a small square kernel.

6. Edges are extracted from the image using Canny Edge Detection, and contours are generated.

7. A bounding box that spans all the contours is found and the height of that box is used as the height of the baby in pixels.

8. Since the size of the blanket is known, both in pixels and in cm, the height of the baby in cm is recovered.

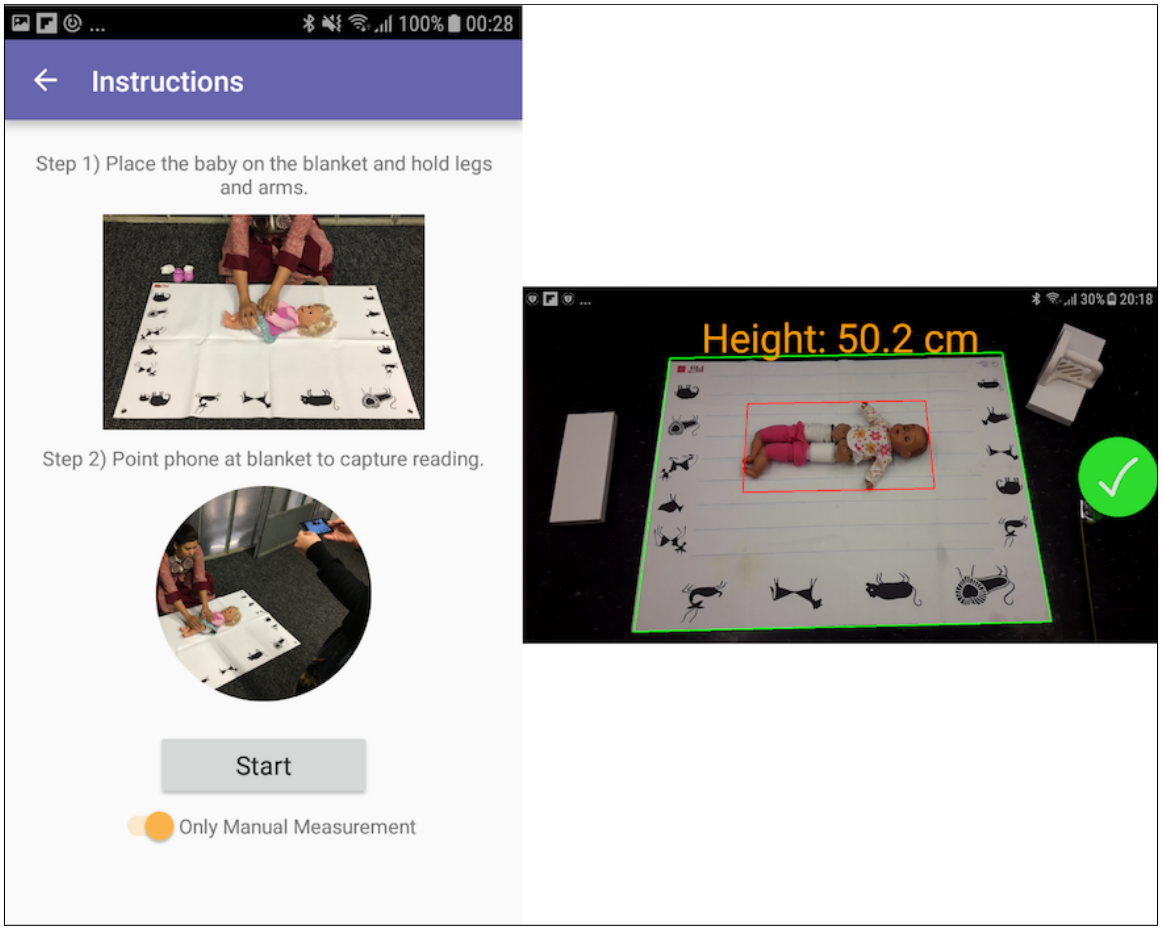

Figure 3-4: Baby Blanket App Instructions Screen (L) and Operation Screen (R)

3.3.3 Length Measurement Lab Validation

It has previously been established that lighting conditions affect the performance of our tools as expected. Work from a previous publication [2] discusses in detail the light intensity and angle requirements for accurate readings. For this section, we assume ideal lab conditions, and compare the automatic app measurement with a manual measurement that was taken as carefully as possible to guarantee the integrity of an accurate length measurement. The setup for the manual experiment is depicted in Figure 3-5.

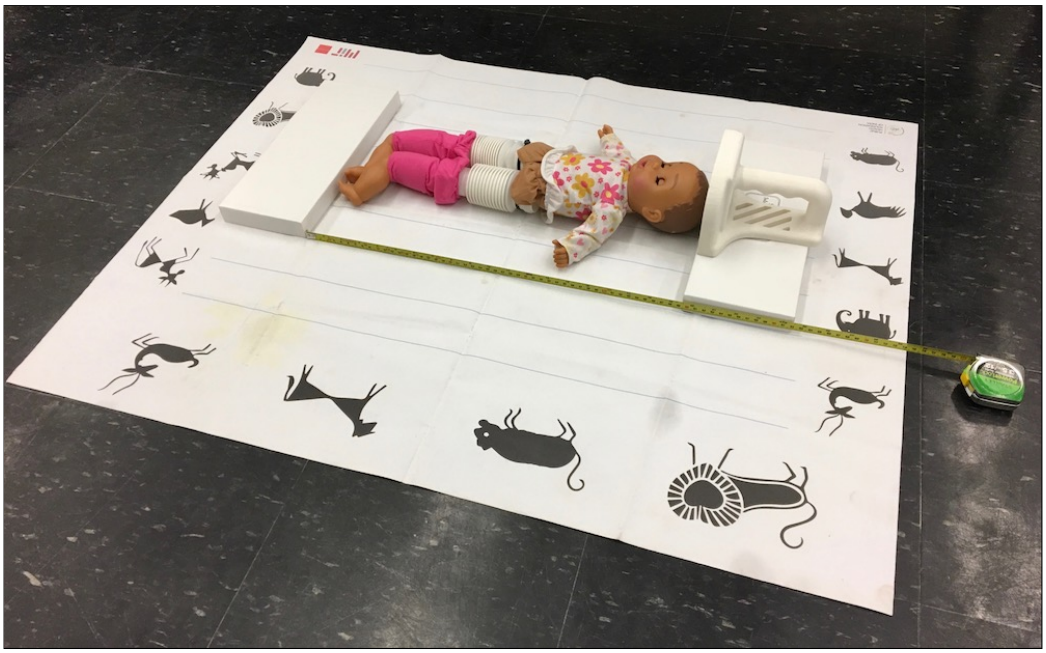

Figure 3-5: Length Experiment Set Up for Manual Measurement

A baby doll was modified with accordion legs that could expand and compress to achieve different heights. The manual measurement was obtained by placing blocks on the feet and the head of the baby doll, and the distance between the edges of the blocks was measured with a measuring tape. For the automatic measurement, the blocks were removed and the length captured using the Blanket app as shown in Figure 3-6. The experiment was repeated for different lengths, starting on 43 cm which is less than 2 standard deviations than the average length of a newborn baby, and went up to 66 cm, a value close to the physical limitations imposed by the size

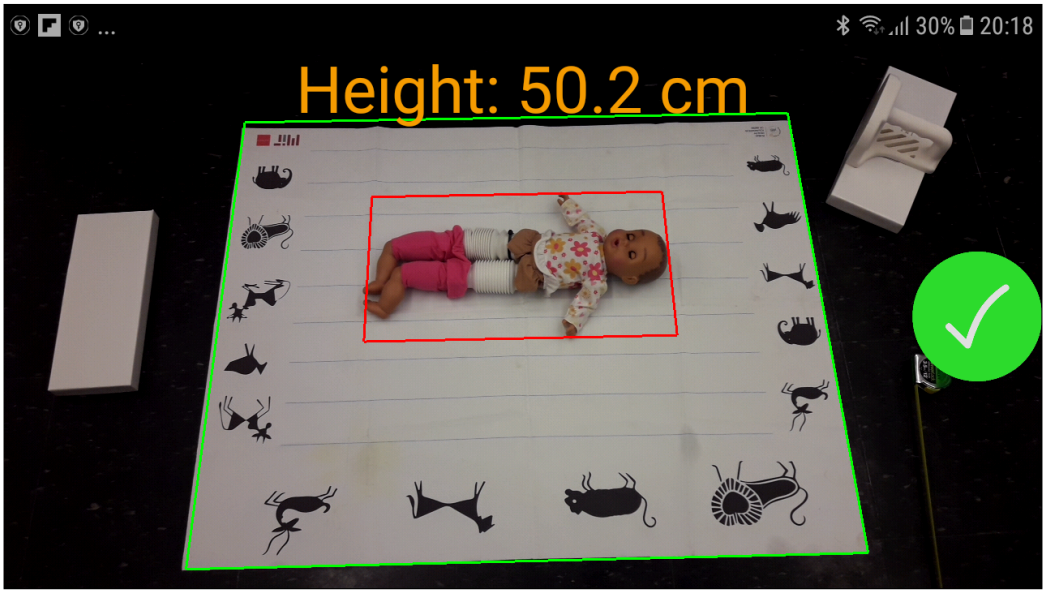

Figure 3-6: Length Experiment Set Up for App Measurement

of our blanket and the restrictions of our algorithm.

The Bland-Altman plot in Figure 3-7 shows a comparison between the manual measurement h_m of the baby's length vs. the app measurement h_a for the range of lengths described above. The mean error was calculated to be 0.91 cm, which suggests the existence of a small positive bias that will be discussed in Section 4.5.1

3.4 Measurement of Weight

3.4.1 Design of AR Target

We augmented the spring scale used by Community Health Workers by placing a cylindrical sticker that spanned the whole circumference of the scale. The image of the sticker was designed manually and started as a friendly tiger face, its stripes and some spots, but then, a factor of randomness was added to all the shapes in the target to generate strong and recognizable features. The height used for the target is 7.7 cm and a circumference of 8.2 cm, which results on an average radius of 2.6 cm. Because our algorithm searches for transitions between the black of the body of the scale and the white from the measuring region and the phone can track the target in

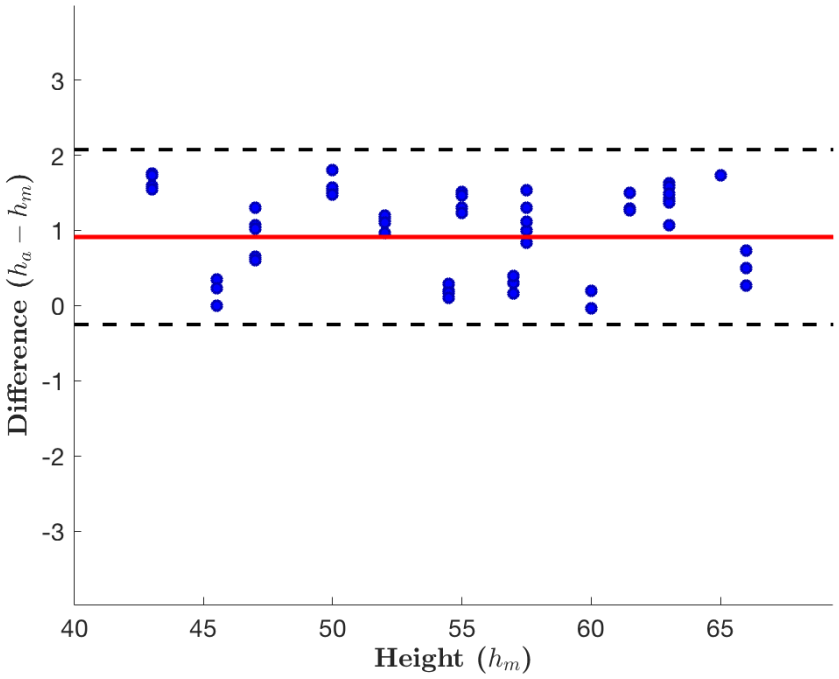

Figure 3-7: Bland-Altman Plot that compares the manual measurement h_m of the baby's length vs. the app measurement h_a for a feasible range of lengths.

360°, an additional white sticker had to be added on the measuring region to span the its whole circumference. The final version cylindrical target and modified scale can be seen in Figures 3-8 and 3-9, respectively.

3.4.2 Weight Algorithm

Finding the tangent plane parallel to camera

In the flat target case, as described in section 3.2.2, there was no need for this step. The algorithm would use the corners of the target to directly rectify the whole camera scene. However, for the case of the cylindrical target, there are no four corners, an invariant-like rectangle had to be generated. Such rectangle is a perpendicular projection of the cylinder into the plane tangent to the cylinder that is parallel the to the camera plane. This rectangle is perpendicular to the plane that minimizes the distance between the camera plane and the cylinder axis (See Figure 3-11). Notice that the dimensions of the rectangle are the height of cylinder, and the diameter of

Figure 3-8: Scale Cylindrical Target flatten out

Figure 3-9: Scale with Cylindrical Target attached

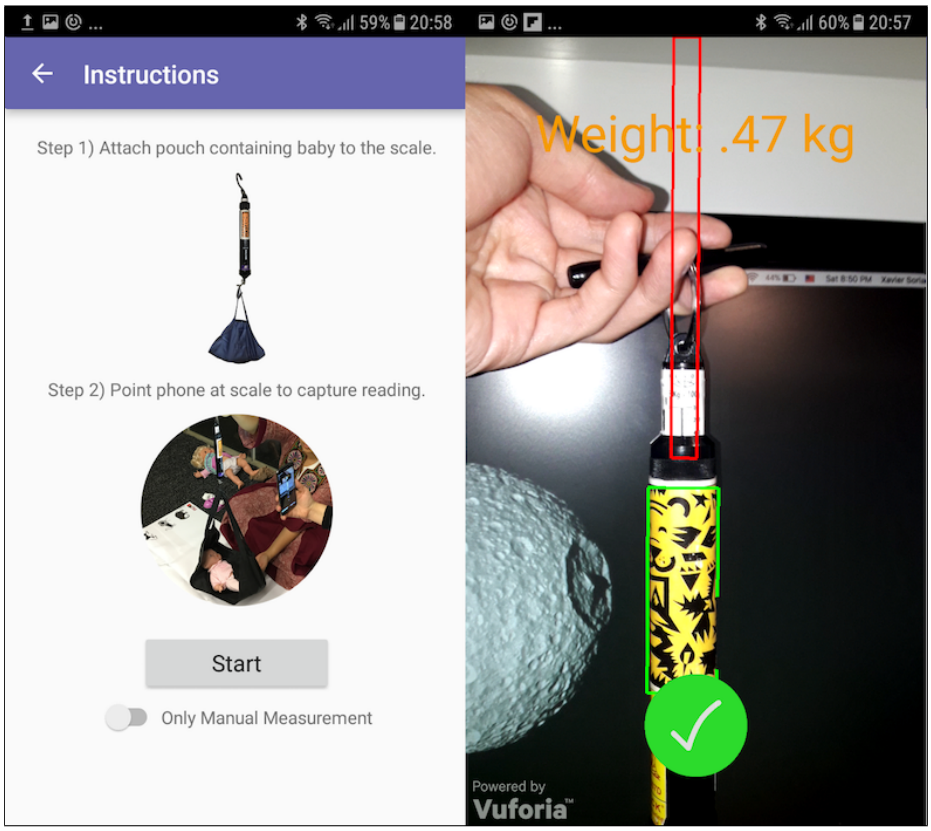

Figure 3-10: Baby Scale App Instructions Screen (L) and Operation Screen (R)

the cylinder.

To keep track of such plane, we used the orientation information provided by the Vuforia API, which is provided as Euler angles with respect to a frame of reference defined by the cylinder of the target. By always compensating for a respective rotation, we are able to keep our rectangle always in sight in front of the camera. And such rectangle gives us the reference to rectify the scene.

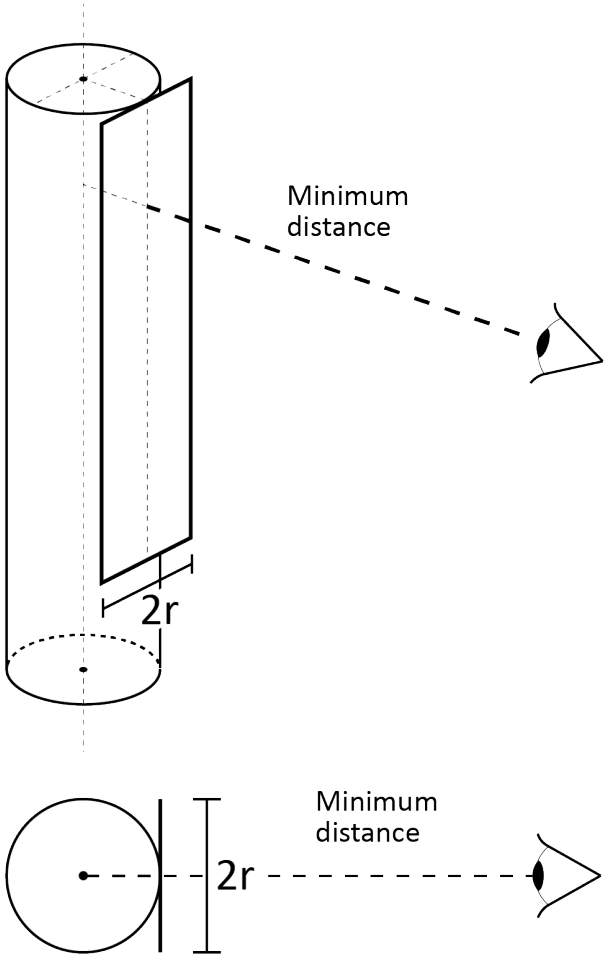

Figure 3-11: Cylinder Tangent Plane Parallel to Camera Plane

Special Rectification, selection of ROI

A rectification step is required to correct for perspective deformations due to tilting of the smartphone with respect to the scale, which is mainly vertical due to gravity. Given the four points in the previous step, we can use OpenCV to find the perspective transform that will turn such quadrilateral into an empirically chosen 700 pixels tall rectangle. Unlike with the blanket, the ROI of the scale measurement is not contained in the target. In this case, a region of interest on top of the target is selected in such way that it spans what a full extended scale would span plus a discretionary buffer region. The offsets and values for the ROI are empirically selected based on the actual dimensions of the scale. The physical height of the white sticker on the scale was 14.1 cm, with white space offset of 1.6 cm. Therefore, a region of interest 16 cm tall and 1 cm wide was selected.

Image Processing of the ROI for Length Estimation

A typical ROI example is displayed in Figure 3-12. The problem of measuring the weight is as simple as differentiating light versus dark regions, and the pipeline required in order to be able to extract the displacement based on such image is described in the following sections.

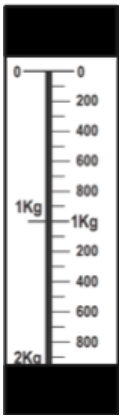

Figure 3-12: A typical ROI from a frame containing the AR Scale

The first challenge is to remove the measuring guides, or black content, in the white region. This can be achieved by dilating using a 1 x 3 pixels vertical kernel and a 7 x

1 pixels horizontal kernel. The horizontal lines that mark discrete measurements of the scale and the thicker vertical line that holds them together are strongly reduced and hopefully eliminated with such operation. However, in order to preserve the vertical information, we erode back using the first 1 x 3 vertical kernel. Since accurate information in the horizontal axis is not needed, we do not apply erosion with the horizontal kernel. This conveniently leaves the frame with an increased white content.

For the sake of removing any noise or remaining black content and working with a smooth region of interest, a bilateral filter [18], which is an edge-preserving smoothing filter, is also applied. With the bilateral filter, we are able to preserve the transitions between dark and light, while at the same time reducing the noise of each region.

The next step is to obtain a 1D array by calculating the average of each row. The assumption is that even if after the preprocessing there is still noise present, the predominant color in the region will turn the average into an accurate representation of that row.

Again, OpenCV is used to perform color thresholding. We first transform the image into HSV, and then the *inRange* method from the *Core* package is used to threshold based on saturation and value. White corresponds to low values of Saturation, and high values of Value. This operation results into a binary array on which 1 means white and 0 means black. A simple 1D edge detector is used to find the transitions between dark to light and light to dark, to get the number of pixels in the light region. If the number of continuous pixels is either black or white and is greater than a certainty threshold, the algorithm determines that the region has transitioned. Given that we know the mapping in between the length of the target from pixels to mm, we can recover how much the scale has displaced in mm as well: 700 pixels map to 12.6 cm, which themselves map to 5 Kg. The calculation is a cascade of simple rules of three.

3.4.3 Weight Measurement Lab Validation

For testing the weighing scale application, the scale was attached to custom made setup shown in Figure 3-13. The setup was made by rearranging pieces a clamp in a

way that it would keep the scale extended at specific weights. These weights ranged from 2Kg, which is about t 3 standard deviations below the weight of an average new born, to 5kg, the upper limit of the scale itself. The comparison between the manually recorded weight and the app captured weight are shown in Figure 3-14.

Unlike the height measurement tool, the weighing tool is used by holding the phone fairly close to the AR target (as visible in Figure 3-10). Therefore, it was possible to use the integrated flash from the smart phone to assist in cases with low illumination. The glare caused by the reflection of the flash on the surface of the target posed no significant complications. Given the correct light conditions, a mean error of 0.0043 Kg was achieved (effectively 0). It is important to notice, however, that the error ranges between $\pm$100 grams, which is non negligible for newborn babies.

Figure 3-13: Scale Experiment Setup for Weight Measurement Validation

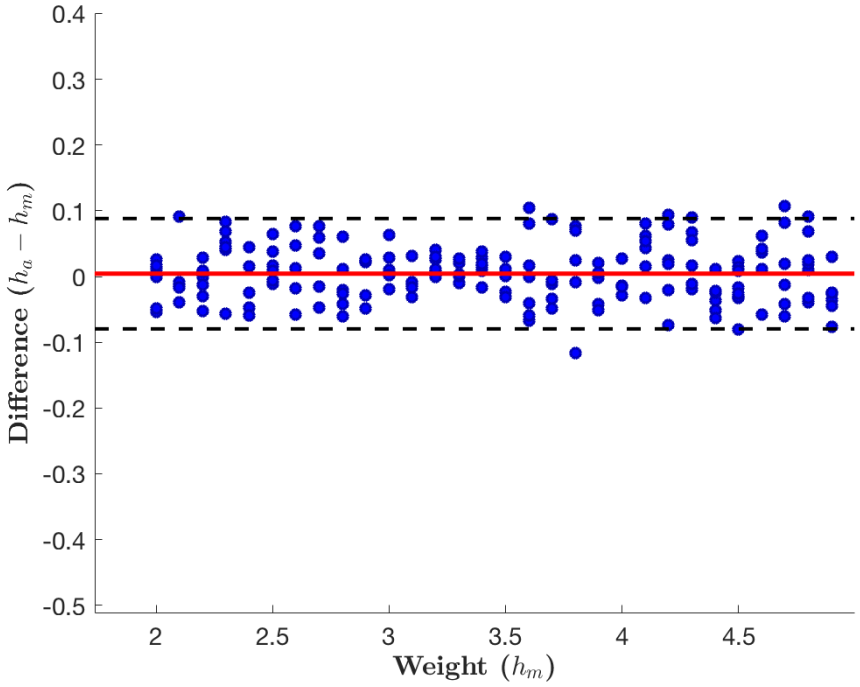

Figure 3-14: Bland-Altman Plot that compares the manual measurement h_m of the baby's weight vs. the app measurement h_a for a feasible range of weights.

3.5 Measurement of Middle Upper Arm Circumference

3.5.1 Design of AR Target

The MUAC measuring tape was modified to include an AR target surrounding the measuring window, i.e. the slot through which the measurement is read. The band was also modified to contain periodic rectangular patterns whose use will be explained in section 3.5.2. The target is placed on the MUAC in such way that the "eyes" occupy a small vertical space, and the "ears" are actually isolated from the main face of the MUAC, and attached only through thin connections. This design choice is meant to minimize the deformation of the target. The dimensions of the target are 2.4×9 cm and the dimensions of the measuring window are 2×1.5 cm.

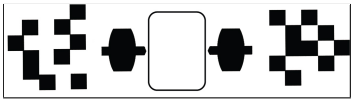

Figure 3-15: Final version of MUAC AR Target

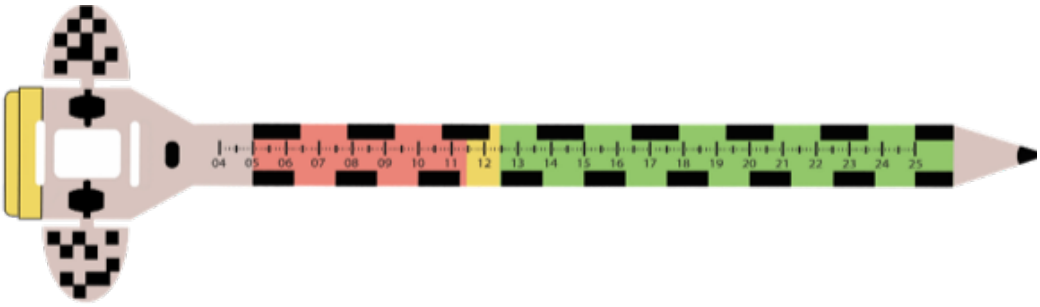

Figure 3-16: Final version of modified MUAC band

3.5.2 MUAC Algorithm

A fundamental part of the algorithm for the MUAC measurement calculation, is the claim that the information contained by the top and bottom patterns that is visible from the measuring window is enough to determine the number in the middle of the window.

Absolute Location based on phase shift

Consider a segment of length 1, that has been divided in equally-sized segments of length T_1 and divided again (Figure 3-18 on the bottom and top, respectively) in equally-sized segments of length T_2, with

$$1 = nT_1 = (n+1)T_2 \tag{3.1}$$

for some $n \in \mathbb{N}$. Let's also define a rising edge of T_i as a transition from white to black when the segment has been divided using the period T_i. We claim that, given a window of the segment, with window length $W > \frac{T_1}{2}$, it is possible to determine the absolute position x of the center of the window, in terms of the distances from x to the next rising edges of T_1 and T_2: L_1 and L_2 respectively and T_1 and T_2 which are, of course, known.

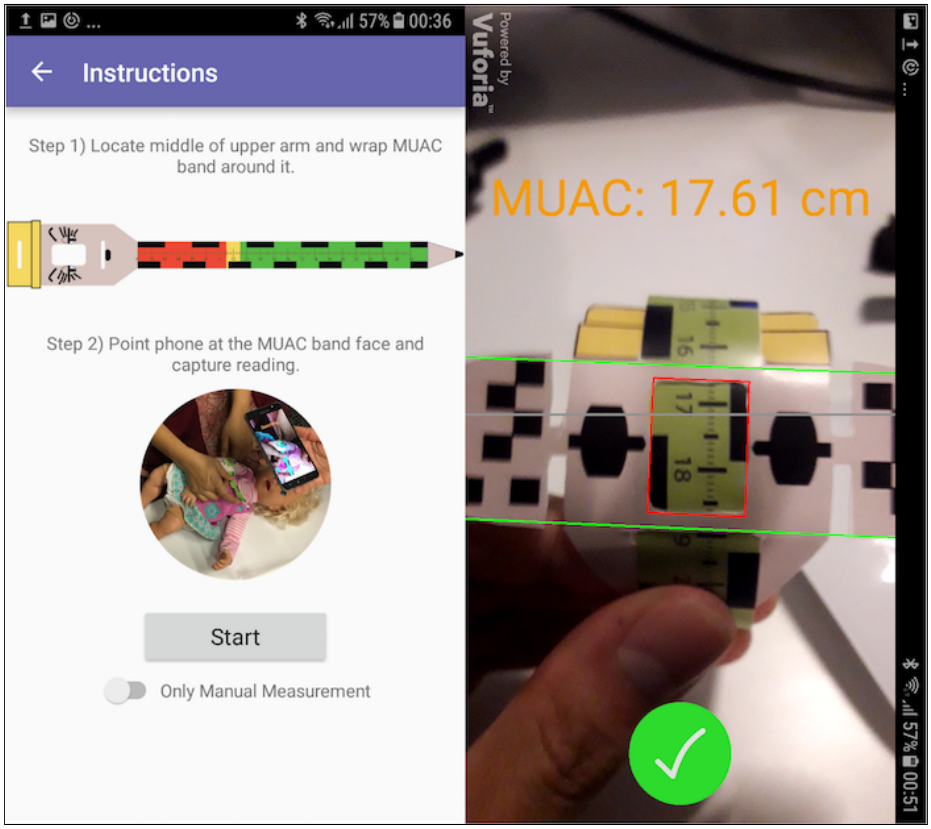

Figure 3-17: Baby MUAC App Instructions Screen (L) and Operation Screen (R)

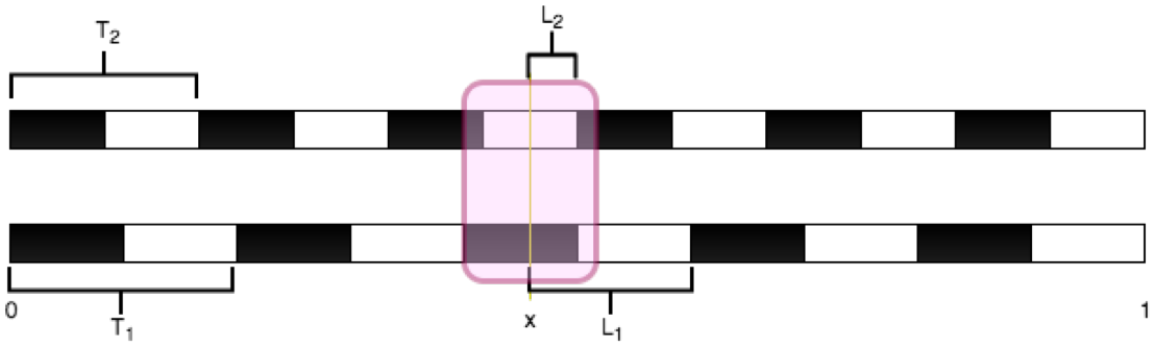

Figure 3-18: Absolute Location based on phase shift — Definition

Notice that we are not guaranteed to get a rising edge in our window for both T_1 and T_2. Nevertheless, $W > \frac{T_1}{2}$, we are guaranteed to get a transition for each of the two divisions. If the transition is from white to black, we are all set: we have a rising edge. However, if the transition is from black to white, we can still infer the position of the next rising edge by adding half of its respective period.

Now, let's observe that there is a pair of positive integers m_1 and m_2, such that x

can be written as

$$x = m_1 T_1 - L_1 = m_2 T_2 - L_2 \tag{3.2}$$

And plugging equation (3.1) into equation (3.2) we obtained

$$\begin{aligned} m_1 T_1 - L_1 &= m_2 \frac{n}{n+1} T_1 - L_2 \\ m_1 - \frac{n}{n+1} m_2 &= \frac{L_1 - L_2}{T_1} \end{aligned} \tag{3.3}$$

Figure 3-19: Phase shift discontinuity schematic

To find the absolute location x we now only need to solve for m_1 or m_2. We will show that

$$m_1 = m_2 - \Delta \qquad \text{for } \Delta \in \{0, 1\} \tag{3.4}$$

In Figure 3-19, the sections for which $\Delta = 1$ have been highlighted in yellow, and $\Delta = 0$ for the rest. Now, observe that $k < n,\ k \in \mathbb{N},\ kT_2 \leq kT_1 < (k+1)T_2$, from which we conclude that equation (3.4) is a trivial consequence.

Finally, we can solve for m_1 based on two cases:

- **Case 1:** $m_1 = m_2$

  First, we have that

  $$m_1 = (n+1) \times \frac{L_1 - L_2}{T_1} \tag{3.5}$$

- **Case 2:** $m_1 = m_2 - 1$

In this case, we have that

$$m_1 - \frac{n}{n+1}(m_1+1) = \frac{L_1 - L_2}{T_1}$$
$$m_1 = (n+1) \times \left(\frac{L_1 - L_2}{T_1}\right) + n \tag{3.6}$$

Practical Implementation

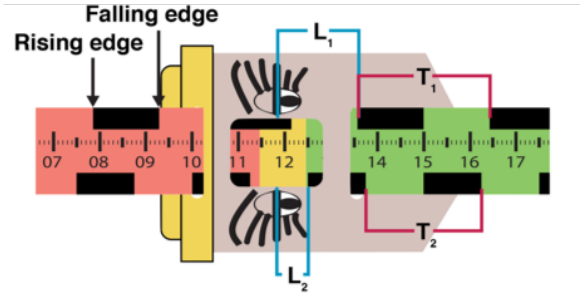

Figure 3-20: Close-up detail of the "head" of an earlier version of the MUAC band which integrates the augmented reality target and optical linear code used to calculate the position along the band.

The formulas from the previous section can be readily applied into the MUAC calculation. As mentioned earlier, each of the long edges of the band contains periodic black markings, which serve as a linear code. One edge has period T_1 and the other edge has period T_2. In our current band design which can measure circumferences from 4 to 25 cm, there are $n_1 = 7$ periods T_1 along one edge and there are $n_2 = 8$ periods T_2 along the other edge. The periods T_1 and T_2 were chosen so that the width of the measurement window would be greater than both $\frac{T_1}{2}$ and $\frac{T_2}{2}$.

In order to calculate circumference, the MUAC software focuses inside the measurement window, and analyzes the phase difference between the black markings on each edge of the band. Just like in the previous section, the value in the center of the measuring window, in this case the circumference, can be found by applying the Equations (3.2), (3.5) and (3.6).

The MUAC software tool (Figure 3-17), extracts a top and a bottom slices from the ROI. For this discussion, the target in Figure 3-15, has been rotated 90° counterclockwise to match the orientation of Figure 3-20. The larger dimension of these

slices is equal to the larger dimension of the measuring window, and the heights (or smaller dimensions) are both equal to a length that empirically spans the height of the rectangular patterns (about 2.5 mm). To find the rising, the algorithm performs a processing pipeline similar to the one explained in detail in section 3.4.2. The slices of the ROI go through a bilateral, column averaging, and color thresholding to end up with a pair of 1D arrays. The process then searches for edges by scanning the image segment starting from the center and moving outwards one pixel at the time and detecting abrupt changes in brightness, which are white-black or black-white transitions in the thresholded image. If the rising edge is found to the left of the center, a full period is added to infer the next rising edge. Alternatively, when a falling edge is found, half a period is added.

At this point, the reader should not be convinced about the correctness of this method, for an important consideration is still missing: the changing curvature of the band. Since the MUAC band is not flat, but rather wrapped around a child's arm, an additional correction must be added to our algorithm to account for curvature. If the radius of the circumference is r and the observed distance is d_p, the corrected distance d_c can be obtained by the formula

$$d_c = r \times \arcsin\left(\frac{d_p}{r}\right)$$

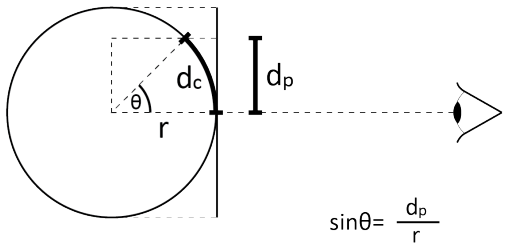

Figure 3-21: MUAC Curvature Correction

However, the radius is not fixed, it is actually dependent on the measurement itself. This issue calls for an iterative algorithm which makes an initial estimate and

enters a feedback loop until converging on a corrected value. Starting with an initial circumference estimate of 15 cm, our correction step uses the measured distances that are projected into the camera plane to correct the circumference value. This new circumference estimate is then passed to the next iteration of the algorithm which continues to iterate until the measurement converges to within 1mm.

3.5.3 MUAC Measurement Lab Validation

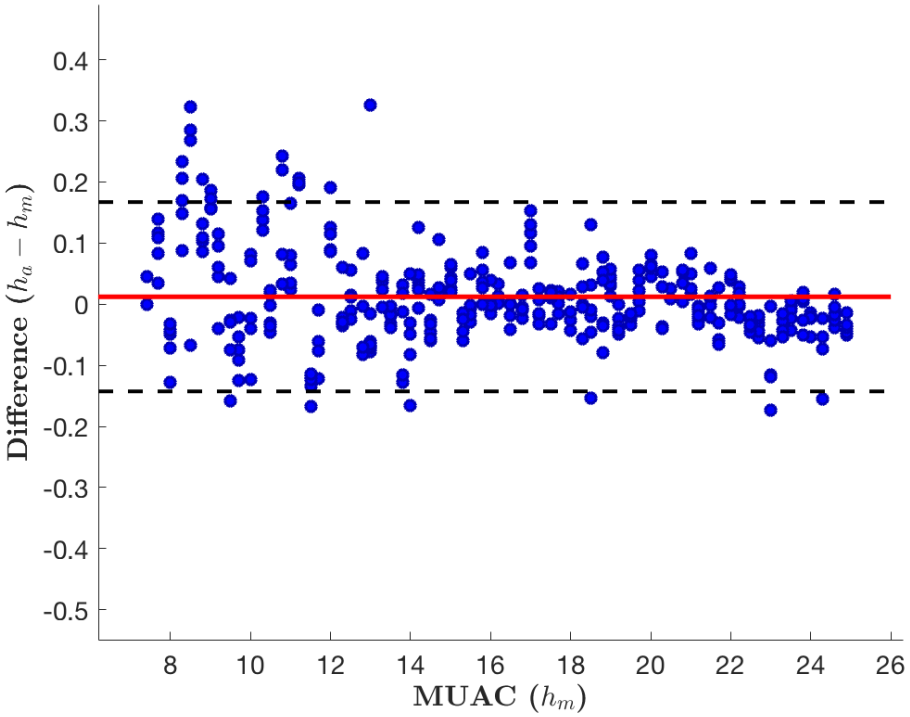

Figure 3-22: Bland-Altman Plot that compares the manual measurement h_m of the baby's MUAC vs. the app measurement h_a for a feasible range of circumferences.

The MUAC measurement tool was tested by manually adjusting and holding the circumference of the band as shown in Figure 3-17(R). The MUAC value measured by the mobile app was compared against the actual value, and the results are shown in Figure 3-22.

Over the entire measurement range (8 cm — 25 cm), good agreement was found between the measured and actual results, with mean error of 0.0122 cm. A few outlier

points were observed in the lower end of the circumference range (< 10 cm), likely due to the increased curvature of the AR target which distorts the geometric calculation.

For this tool, the integrated flash was turned always off to avoid specular reflections and glare that interfered with the measurement. Hence, as with the blanket measurement, the lighting in the scene must be bright enough to allow for a good distinction between dark and light zones in the band.

Chapter 4

Community Health Worker Field Study

Figure 4-1: ASHA workers during the field study

A field study was developed in conjunction with the Public Health Foundation of India (PHFI), which implements much of the training for ASHA workers. Since the ASHA health worker program is a government program, we obtained approval from the Delhi State Health Mission to conduct a pilot study with 13 ASHA workers in

the New Delhi slums using our technology. The study took place between the months of April and September of 2017, on which the ASHA workers were introduced to our kit of mHealth tools and were asked to include the use of our kit in addition to their normal routine.

This study was sponsored by the generous support of the TATA Trusts and the MIT Tata Center.

4.1 Motivation and Study Aims

With the rapid growth of mHealth tools in the developing world, it seems like the natural step to put our tools in the hands of CHWs right away. However, before any large deployment it was important to show the tools as a proof of concept to establish the feasibility of their potential. The original sole and main aim of the study was measure the user acceptance of our kit within ASHA workers and also to explore the feasibility of using mHealth tools as a job aid for ASHAs and its field validation. However, as early as the first days of training, the high variability of the scenes during operation proved the study to be a very adequate setting for validation of our tools and algorithms in real life field conditions.

The study also served as means to observe how ASHAs adjusted to mHealth Apps at their own pace in their daily work flows, so that influential factors in the adjustment could be taken into consideration for a more sustainable scale-up model.

We can summarize the objectives of the study as follows:

- To assess the feasibility and the ease of use of MHTs as a job aid for ASHAs as they perform their daily duties in their usual resource-constrained settings.
- To validate and improve our algorithms and the design of our hardware, e.g. AR targets, blanket material, scale's sticker, MUAC material.
- To visualize, understand and relate the geographic and demographic distribution of the families of the study back to the results from the assessment from our kit.

4.2 Mobile App Design: Baby Naaapp

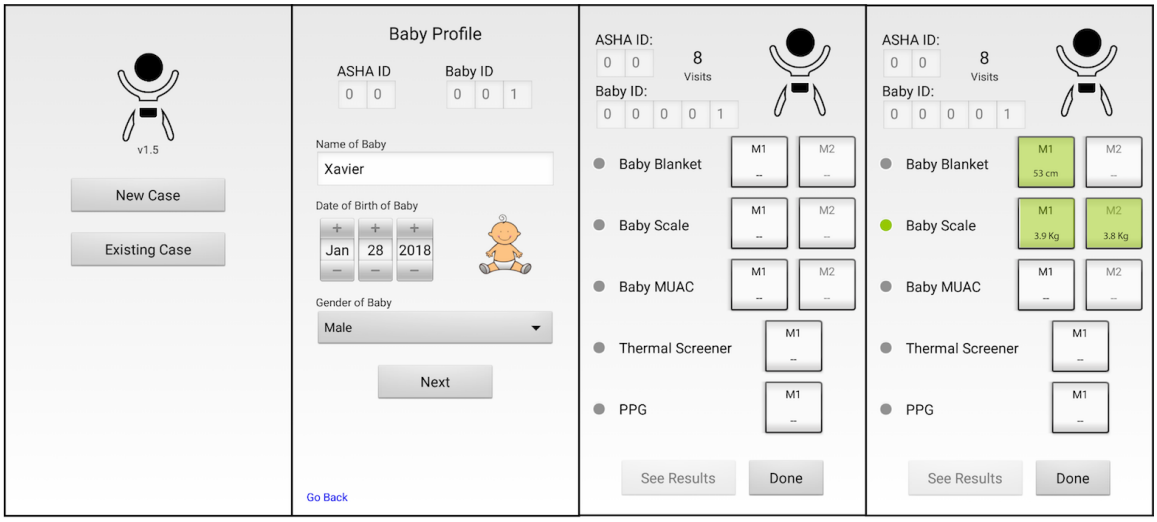

Figure 4-2: Screenshots of Baby Naapp Workflow

To facilitate data collection during the study, a master app was developed and was given the name Baby Naapp, as an allusion to the Hindi word "naap" which means "measurement". The purpose of Baby Naapp was to

- Create a centralized hub to easily create and retrieve baby and family profiles
- Provide a checklist and summary of the measurements that ASHA workers had to perform, allowing them to call each individual measurement from Baby Naapp itself
- Compile and organize the data from each measurement in convenient way for analysis, while at the same time appending information about timings of the usage of the tools and GPS coordinates.

The app was developed on Android 6, using a Samsung Galaxy J7 (J710MN) from 2016. For this reason, the J7 was our phone of choice for the study.

4.2.1 Personal Profiles + Demographic Data

With Baby Naapp we introduce the concept of profiles that can be shared by different apps in our lab. The details of the implementation of such framework can be found

in section 6.2. The profile questions were divided into two groups: first, the basic profile questions, which contained the child and mother information — and second, the additional profile questions, which would provide us with insightful information about the socioeconomic status of the baby and its family.

Only the Basic Profile 1 Questions were mandatory as they where necessary to perform and analyze the measurements. The questions from the Basic Profile were simple enough for the ASHA workers to understand them in English. However, for the more detailed additional questions, translations in Hindi were presented. The list of all the questions can be found in Table 4.1 and the screens that contained the translated screens can be seen on Figure 4-3

Basic Profile 1: Baby Info	**Additional Questions: Socio-Economic Info**
Name of the Baby _____ **Date of Birth of Baby** ___ ___ _____ (DD MM YYYY) **Gender of Baby** □ Male □ Female	**What type of ration card does your family have?** □ National Food Security Card □ Above poverty line (APL) □ Family does not have any ration card **In the past 2 years, has a doctor ever told you that you have Anemia?** o Yes o No
Basic Profile 2: Mother Info	**Type of food your family eats: (you may check more than one box)** □ vegetarian □ non-veg
Name of Mother _____ **Date of Birth of Mother** ___ ___ _____ (DD MM YYYY) **Address** _____ **Phone Number** _____	**What do you use for cooking? (you may check more than one box)** □ wood stove □ charcoal □ LPG stove

Table 4.1: Profile Questions

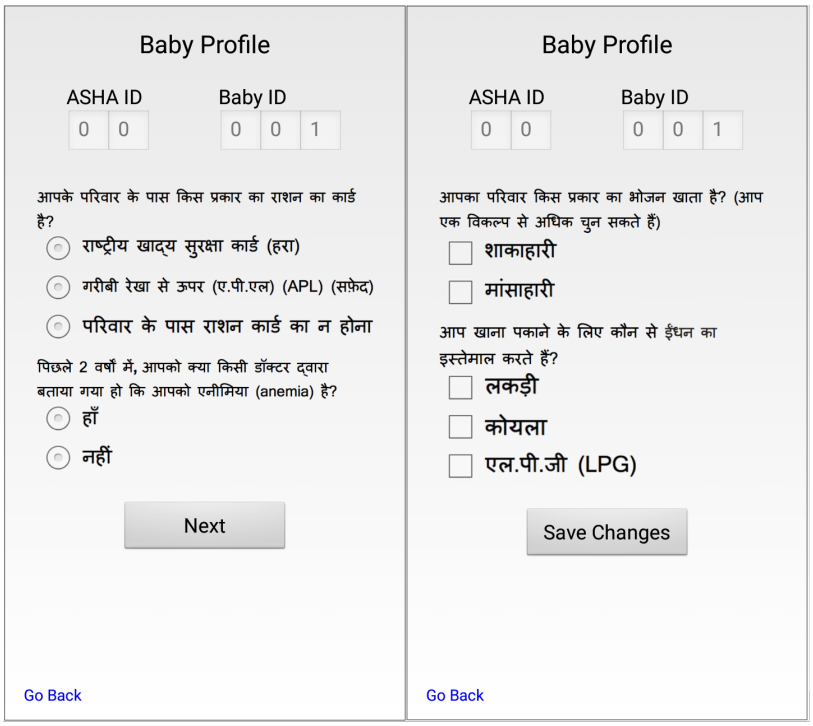

Figure 4-3: Screenshots of Additional Profile Questions in Hindi

4.2.2 Anthropometric Measurements

The AR Anthropometric Measurement apps that were described in detail in chapter 3 were on of the most crucial components of the study. For the Scale (Weight) and MUAC apps, an additional screen was added to enter the measurement that the ASHA would normally record in the absence of mHealth tools. Their manual measurement served as a comparison reference for the measurements automatically recorded by our tools. For the Blanket (Length) app, there was no comparison, since ASHA workers generally do not carry interferometers to perform a manual measurement.

The summary screen on Baby Naapp provided ASHAs with a list of the measurements included in this study. For the anthropometric measurements, two buttons are present next to the name of the tool. When pressing the first button, Baby Naapp automatically passes the baby information to the measurement app, and starts the measurement activity. Upon the completion of the app and manual measurements, the measurement app returns to Baby Naapp. The button is then highlighted with a green tint and the value of the manual measurement (or the app measurement if the

manual is not present) is written on its body. The second button is enabled and the ASHA worker repeats the process for a second time. This interaction result is shown in the two leftmost screenshots of Figure 4-2.

In order to avoid biased data in the results, the ASHAs were never presented with the value of the automatic measurement. Instead, they are only informed that the measurement was ready to be saved by a "Ready" message on the top of the screen, as well as the operation button changing from a grayed-out X to a green tinted check mark.

The ASHA workers were instructed to perform the length and weight measurements on new born babies only (younger than 6 months old), whereas the MUAC measurement made sense only for older children (2 up to 5 years-old).

4.2.3 Physiological Measurements

The study also served as the perfect environment to introduce early stage physiological tools that could be used in the future to screen for disease or abnormalities. While the measurements performed by these tools were no overly complex, they granted us important impressions on how ASHA workers and families would adapt to these more "advanced" tools.

Thermal Assessment

A custom mobile app was created that enables the health worker to quickly capture a thermal image of the infant that is lying in supine position. We made use of a commercial USB thermal imaging module designed for use with mobile phones. For guiding the health worker, a silhouette of a child's body is superimposed on the thermal imaging screen, and the health worker was instructed to make sure that the baby's head is aligned with the silhouette. This helps ensure that the infant will be centered in the image and that the distance to the infant will be consistent (Figure 4-4). Applications of thermal analysis in global health and a more advance thermal tool are explored in sections 5.3 and 5.4

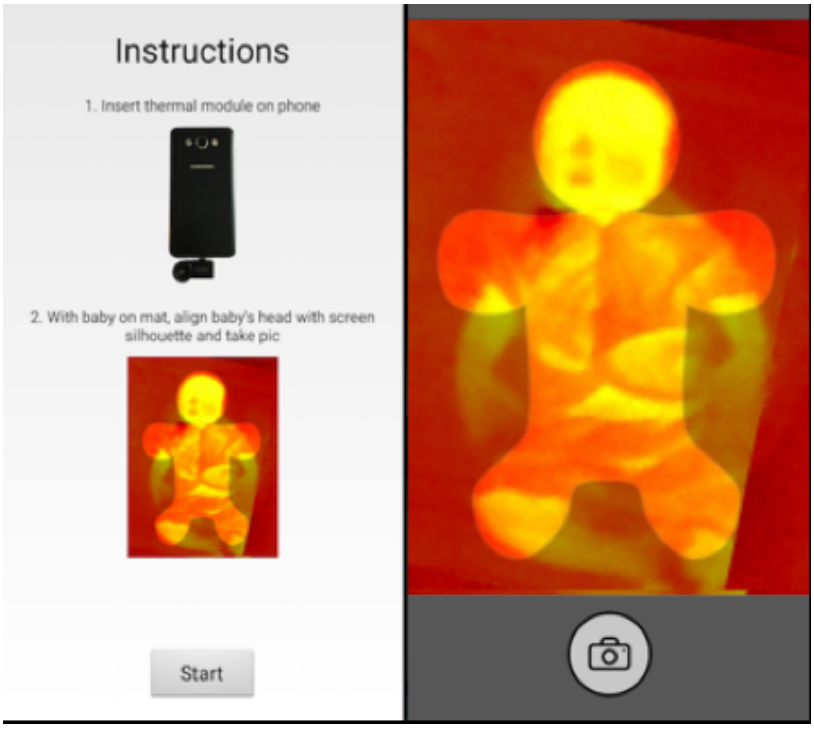

Figure 4-4: Baby Thermal App Instructions Screen (L) and Operation Screen displaying how the head of the baby has been aligned with the head of the baby silhouette (R).

Cardiac Assessment

Figure 4-5: Photo of custom PPG device connected to Android phone (L). PPG device being applied to baby foot in a field study (R).

Cardiac assessment is fundamentally important for newborns and infants. Pho-

toplethysmography (PPG) and pulse oximetry are standard tools for pediatric care [19] and the value of PPG devices has been demonstrated in global health [20] for congenital abnormalities [21] and screening for sepsis and infectious diseases [22].

We developed a custom PPG device (Figure 4-5), for use with Android phones, which contains 3 different illumination wavelengths and is capable of measuring both reflected and transmitted light from a baby's foot. The device was developed as part of a mechanism to measure arterial stiffness using pulse wave velocity (PWV). Three PPG clip-on probes were used, one on the toe, one on the finger, and one on the ear. These devices do not require batteries and plug into the USB port on the phone. For the ASHA study, an adaptable custom probe was also designed so it could be wrapped around a neonate's foot as well as an older child's finger. While this tool has many potential applications, for the purpose of the current study, we have initially implemented only heart rate and heart rate variability for field use. The design, manufacturing, operation, and analysis of the PPG device are examined in detail in [23].

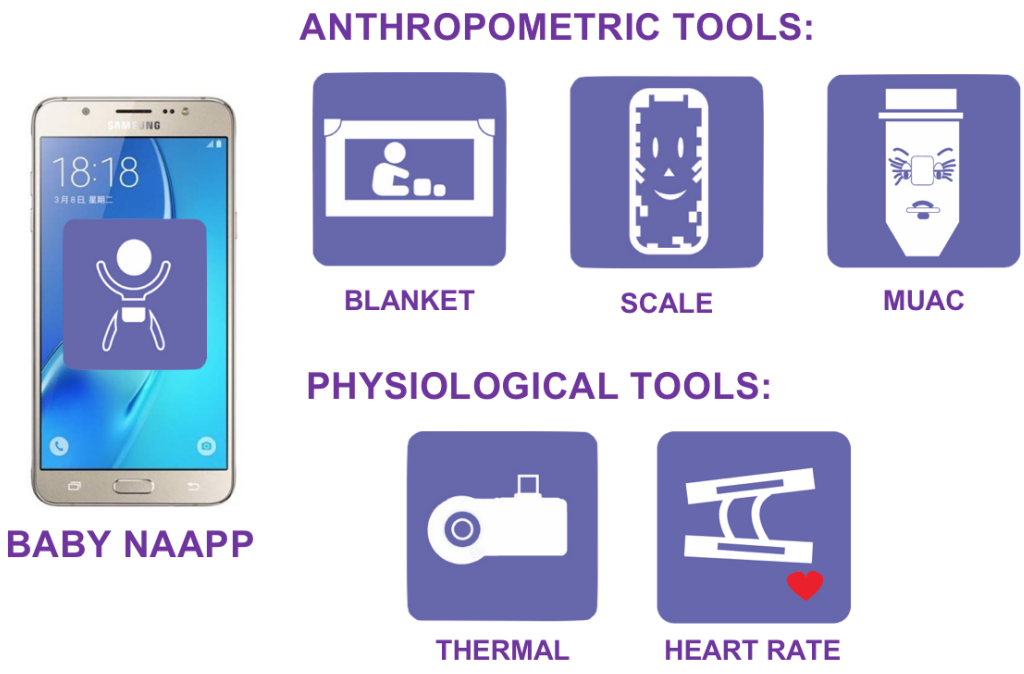

Figure 4-6: Baby Naapp and the ASHA Kit Icons

4.3 Field Study Design and Methods

4.3.1 Participants and Location

The subjects of the study were 13 ASHA workers whose ages ranged between 34 and 52 years. They covered a total population of approximately 28000 people around different areas of were working in different areas of Juanapur Village, which is situated in the South District of Delhi. Their educational qualifications ranged from seventh grade to graduation and their work experience as CHWs ranged between 2 and 11 years. The majority of the ASHAs did not have any prior experience of using smartphones, but they were all enthusiastic to participate in the adoption of mHealth tools. For their participation in the study, all the ASHAs were remunerated via incentives that were arranged by PHFI.

4.3.2 Training and Data Collection

Training at the Juanapur Delhi Government Dispensary

Figure 4-7: Picture Dr. Suparna leading the training Session Day 1

The 13 ASHA workers received a two day training on how to properly use Baby Naapp and each of the mHealth tools. It took place on the 22nd and 23rd of May,

2017 at the Delhi Government Dispensary in Juanapur and it served to kick-start the study. It was led by doctor Suparna Ghosh-Jerath, a nutrition expert from the PHFI and Indian Institute of Public Health and the principal investigator on the Indian team.

The first session served as a reinforcement on the importance of the CHWs work to tackle the problems of mother and children's health and as a reminder and retraining on the proper way of collecting their routine measurements. The ASHAs also received a comprehensive training on the smartphone, its components and its software. For the last part of the first day, Baby Naapp and the anthropometric tools were introduced and their relationship pointed back to the undernourishment assessment. A short hands-on session was conducted at the end on which the ASHAs were instructed on how to use the Blanket and Scale apps to measure length and weight.

The second session was mostly hands-on, devoting the whole day on teaching the ASHAs the proper techniques to use the tools, such as the correct alignment of the babies in the blanket and tricks to help the algorithms converge the correct answer faster. As mentioned earlier, the ASHA were not discouraged by the slightly steep learning curve of introducing mHealth tools in their routine. Instead, they were rather excited and eager to learn and adapt to the use of new technologies.

Development of study and Data Collection

From the end of May until the end of September 2017, the ASHAs included the mHealth tools in their routine. Each of them was given a Samsung Galaxy J7, which contained a pre-selected hard-coded ASHA ID from 1 to 13. They all used the same phone as a way of reducing difficulty of learning different user interfaces, and as a way of guaranteeing that measurements were collected with a standardized tool.

During the first month, a handholding period was held. During such period, the ASHA workers made their visits in the company of two researchers from PHFI, who mentored them and and assisted them in learning the proper technique and correcting any errors while using our mHealth tools. As some ASHAs grew more confident and proficient on the tools, the PHFI researchers used their judgment to dedicate their

Figure 4-8: Pictures of Hands-On Training Session — On the left, the ASHAs are learning how to use the blanket tool and on the right, the ASHA are being shown how to use the MUAC tool.

Figure 4-9: Pictures of Visit with ASHA Worker — On the left the ASHA is taking the weight measurement, and on the right she is taking the length measurement.

time and resources with the ASHAs who were slower learners.

The PHFI researchers also stayed in constant contact with the technical team and based on their observations and conversations with the ASHA workers, valuable feedback was incorporated into the updates of the different tools. The timeline and nature of this updates is discussed in section 4.4.2.

Data Handling and Aggregation

Figure 4-10: Picture of hotspot set up used to upload data to a remote server

In general, ASHA were not given a SIM card, and were asked not to use one. While one of the goals of the study was to allow ASHAs to get familiarized with the smartphone technology, it was also important to prevent the integrity of the collected information and the resources of the phone from viruses and spam data. The ASHAs were asked to inform the team of any apps that they would like to download, and they were encouraged to tinker with different apps in the phone.

For this reason, no cellular data was used. Instead, two battery powered Wi-Fi hotspots were set up at the government dispensary, with 8-hour battery life –âĂŞ and they were automatically charged using a simple mechanical timer. Whenever the ASHA workers returned to the government dispensary, they were ask to push a button so that the data would be automatically uploaded to the remote server.

4.4 Study Results

4.4.1 Demographic Survey Results

As expected, the population that the ASHAs visited was mainly comprised of families in the lower end of the socioeconomic status. They measured 761 different babies in 1066 visits . Of these babies, 398 were boys and 363 were girls. Interestingly, 524

babies were visited or measured only once.

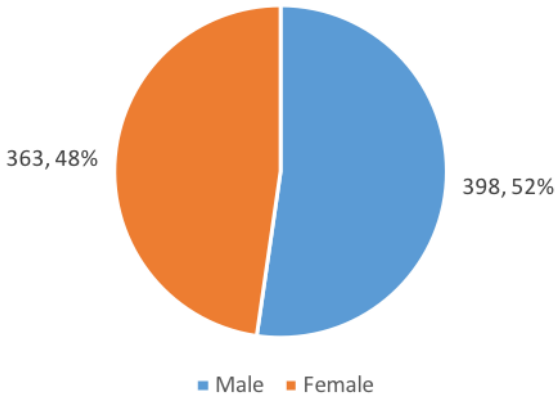

Figure 4-11: Pie Chart of Gender of Babies

Mean	90.31	**Younger than 10**	229
Median	74.60	**Between 10 and 52**	94
Quartile 1	3	**Between 52 and 156**	233
Quartile 3	160	**Older than 156**	205

Table 4.2: Distribution of Babies' ages in Weeks

The distribution of the ages of the babies during their first visit, as shown on Table 4.2 and Figure 4-12, has the particularity that a large quantity of babies were measured before they turned 10 weeks old. This feature is to be expected, as the ASHA manual [24] instructs them to provide close and frequent follow-ups to newborns on the days 3, 7,14, 21, 28 and 42 after the baby is born (also on the 1st day for home delivery cases).

The mother info also shows a skew on the lower end of the age range. While 342 profiles were not filled up (or in some cases filled incorrectly, e.g. 5 year-old mothers), 43% of the respondents mothers were younger than 25 years-old. The distribution of the ages can be seen in Table 4.3.

The results of the additional profile questions agreed with expectations of the socioeconomic information from families in the lower socioeconomic status. 59% of mothers had been diagnosed with anemia within the 2 years before the study, 68% of families were below the poverty line and had a National Food Security Card, 34%

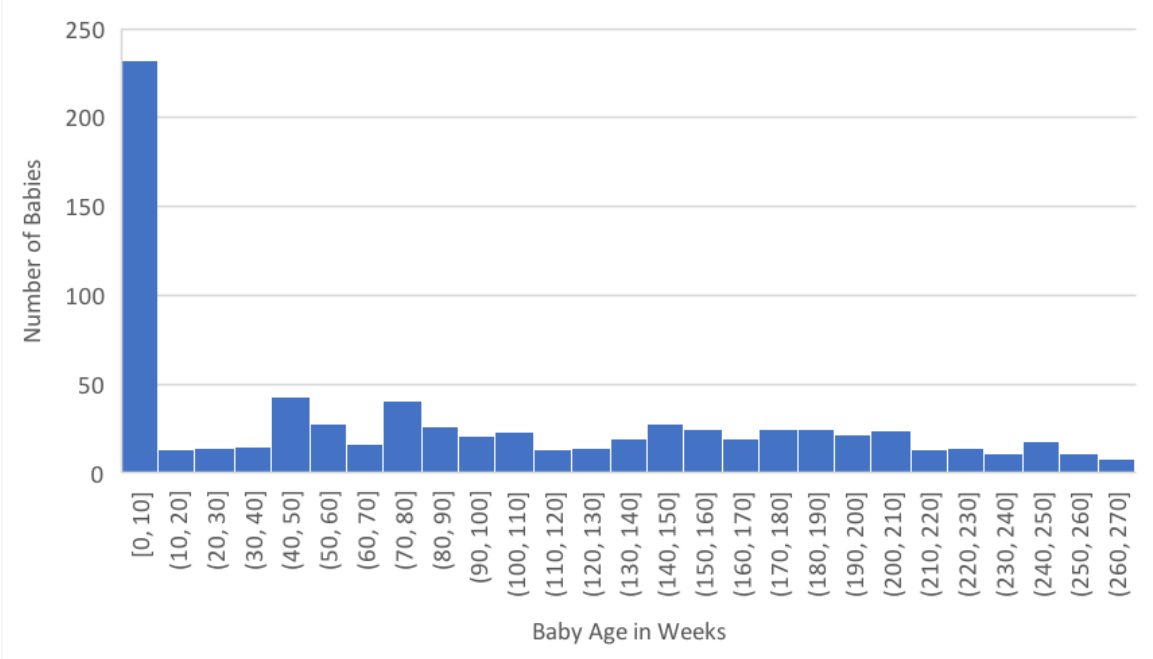

Figure 4-12: Histogram of Babies' Ages in Weeks

Mean	26.37	**Younger than 25**	181
Median	25.65	**Between 25 and 30**	157
Quartile 1	23.02	**Between 20 and 35**	60
Quartile 3	29	**Older than 35**	21
		No Info	342

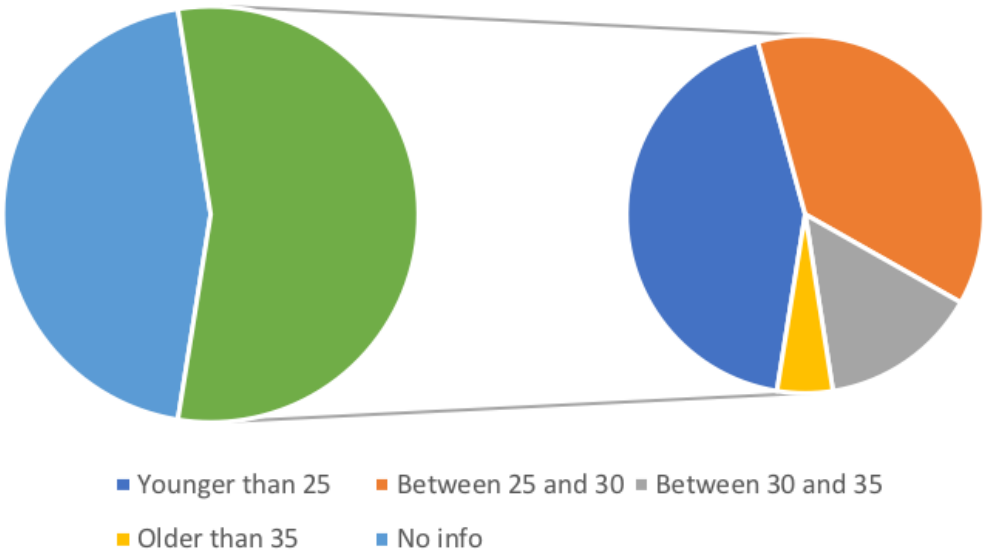

Table 4.3: Distribution of Mothers' Ages in Years

Ration Card		Anemia in Last 2 Years	
National Food Security Card	341	Anemia Positive	312
Above Poverty Line	83	Anemia Negative	216
No Ration Card	80	No Info	233
No Info	257	**Cooking Supplies**	
Food Type		Only Wood	316
Vegetarian Only	184	Only LPG	45
Non-Vegetarian Only	253	Only Charcoal and LPG	6
Both	101	Only Wood and LPG	161
No Info	223	All	7
		No Info	226

Table 4.4: Results of the Additional Profile Questions

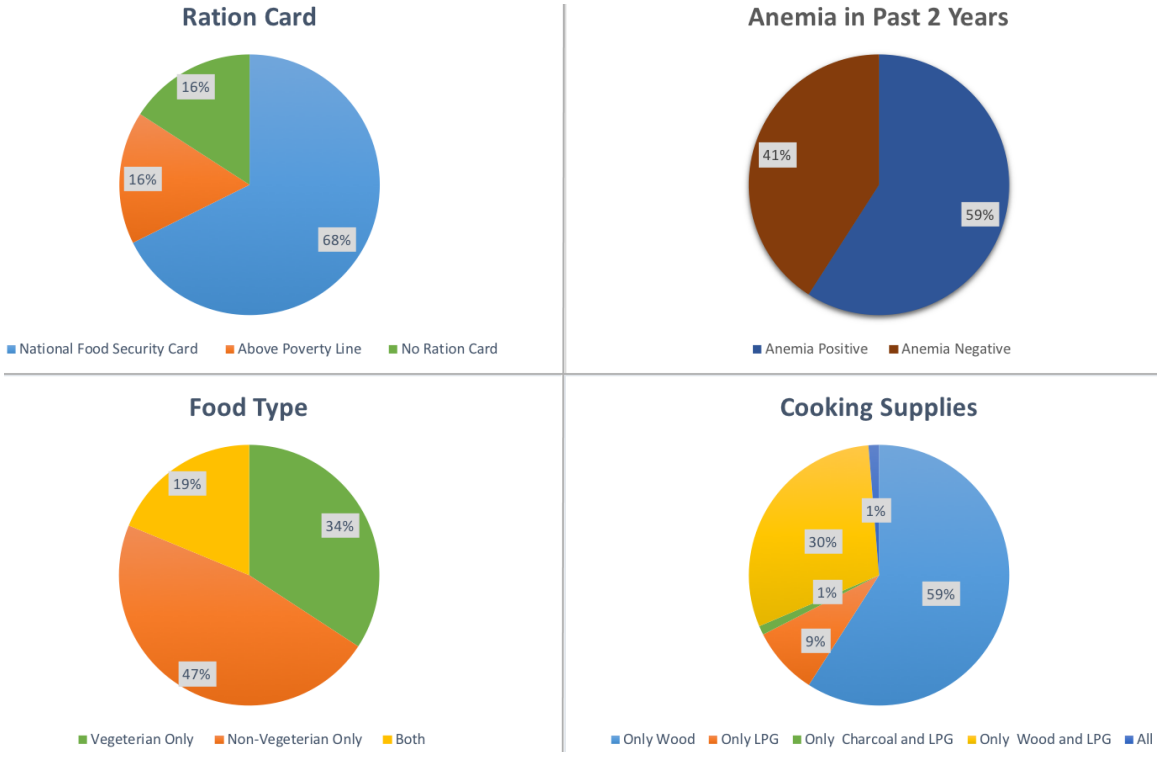

Figure 4-13: Pie Charts of the Additional Profile Results

of the households were exclusively vegetarian and 68% were cooking with biomass stove in their houses. The details behind these percentages are shown in Table 4.4

and Figure 4-13.

4.4.2 Software Updates and Design Iterations

During the the duration of the study, the technical team responded promptly to the complications and challenges that were observed during the visits. In this way, updates were pushed accordingly. In this section we explore the motivation behind and the timeline of such updates.

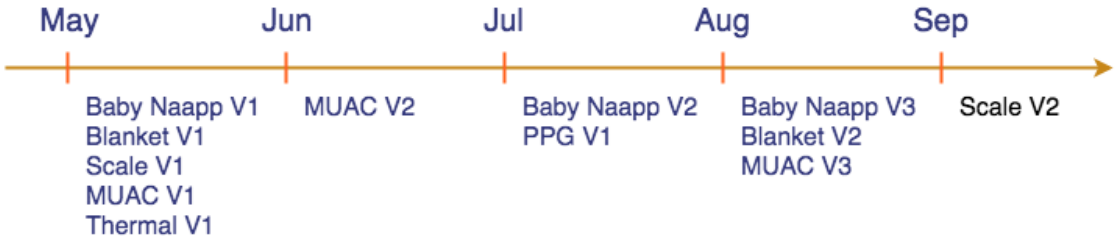

Figure 4-14: Timeline for introduction of updates

Length Measurement

In the initial stages of design, the AR target for the blanket was very simple and uninspiring. In the spirit of creating tool that was more relatable to mother and children, an Indian animal-inspired design was implemented. By this choice, not only the aesthetic aspect of the blanket was improved, but because of the additional randomness in the animal shapes, also the quality of the AR target was significantly increased. The evolution of the target can be observed in Figure 4-15.

During the study, the first problem that arised was that the ASHAs and the mothers were having trouble understanding the correct alignment of the baby in the blanket. Secondly, the space in the houses was very limited and at times the ASHA was not able to find a flat area large enough to extend the blanket. For these reason blue horizontal guide lines were added to the blanket, and its size was reduced to the dimensions presented in section 3.3.1. While such size reduction represented also a reduction in the range of ages the could be measured, it allowed for the ASHAs to meet the space constraints of many households.

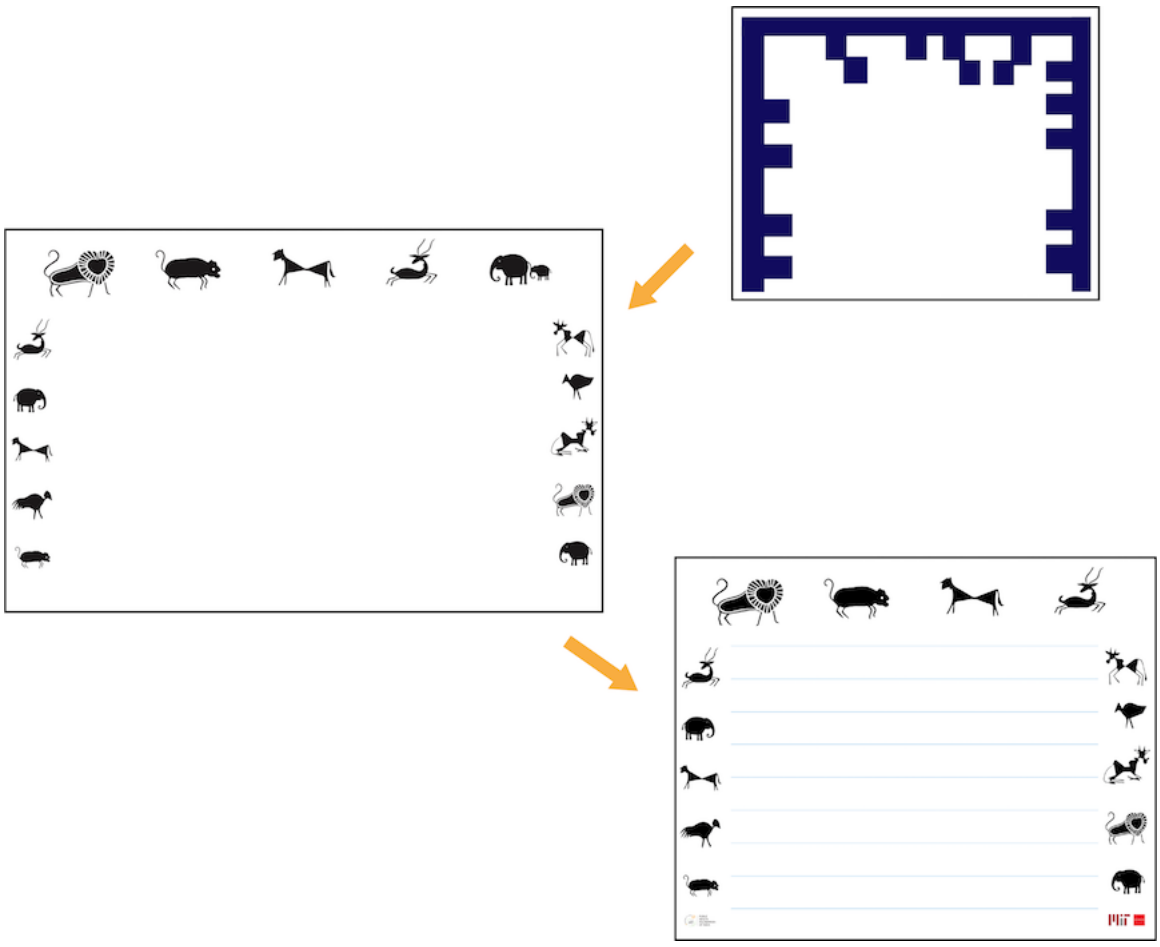

Figure 4-15: Evolution of the Blanket Tool

Weight Measurement

The Scale was the most challenging tool for the ASHAs and the one that took us the most time to correct accordingly. The original design of the scale AR target was a flat sticker, a flat image target. Also animal inspired, it started as a small single tiger face. However, because the shape scale is cylindrical, we tried to overcome the incongruence by creating a target that was long and thin, extending the tiger face and adding a little bit of its body. Nevertheless, the curvature of the target caused tracking to be unstable and difficult to achieve.

During lab testing, this fact was not a terrible problem, but for the ASHAs in the field it was the most frustrating tool. For it to function properly, they had to focus the camera very close to the sticker and slowly move the phone back to be able to

capture the whole frame in the phone screen. This process added to the inexperience of the mother as an aid, cause extended delays and discomfort for the ASHAs and the mothers who had to stay in uncomfortable physical positions while trying to make the measurement converge. The resources were exhausted for the flat target, since increasing its height would cause the ASHAs to have to bring the phone even farther away to capture the reading, while increasing the overall size of the target would generate more problems to preserve the automatic tracking.

A different approach had to be taken and we started exploring the possibility of using an cylindrical target which would indeed be more appropriate for the scale. The cylindrical target extended the flat target by adding an additional tiger on the diametrically opposite side of the first one and by increasing the randomness of the tiger markings. With such change, the tracking became extremely reliable, but a white sticker had to be added around the measuring region so that measurements could be capture in 360° along the axis of the scale.

The evolution of the scale AR target can be seen on Figure 4-16 and the last two versions of the scale can be seen on Figure 4-17.

Figure 4-16: Evolution of the AR Sticker for Scale

Figure 4-17: Final two stages of the Scale Tool

MUAC Measurement

The nature of the MUAC and its not fixed curvature was also problematic. Because of this, we tried create a target that spanned in the direction of no curvature and minimized the content in the direction on which the band curves. The original AR design was a pair of eyes, eyelashes and eyebrows. To accommodate this target, a large MUAC head was included, but this addition came with usability problems. The larger head was difficult to wrap around children's arm. Therefore, an update was made in which the head of the MUAC was reduced.

On the software side, we included first some semi-transparent images of the eyes AR target, that stayed constant on the screen of the phone when the camera was tracking. The idea behind these eyes was to provide guidance on the optimal distance and orientation to track the AR target and get the measurement. However, the natural tendency of all users was to try to perfectly align the super-imposed eyes with the physical target. This interaction defies the purpose of using automatic tracking,

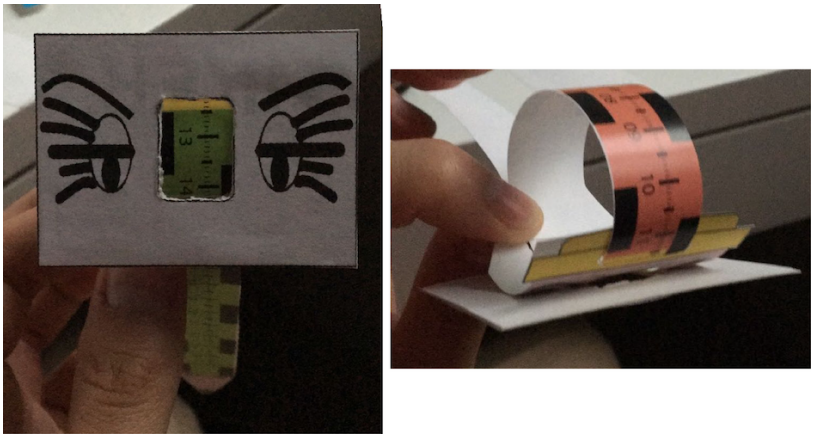

Figure 4-18: Pictures of MUAC Design Attempt with Flat Target

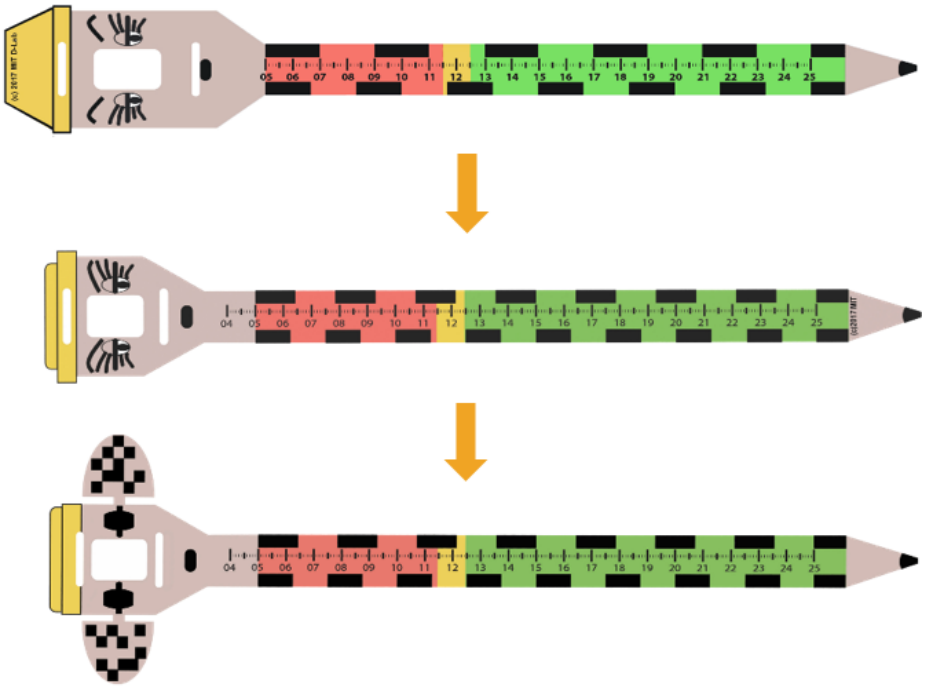

Figure 4-19: Evolution of the MUAC Tool

by going back to putting work on the hands of the user for alignment, and therefore the semi-transparent feature was discarded in the first update as well.

In an attempt reduce the effects of the curvature problem, a couple of designs were explore. The first attempt, shown in Figure 4-18, was to completely eliminate the curvature problem by appending a piece of stiff cardboard that would contain the target, and would be attached to the band only by a thin line. In this way, the target

size was no longer constrained and the band could freely wrap without deforming the target. The solution was very attractive, but not very practical: the cardboard casted shadows on the band that were interfering with the measurement. Since the camera illumination would also interfere with the measurement by causing unwanted glare, this design was discarded and never put in the field.

The final version of the MUAC one was the one shown earlier in Figure 3-16 and can also be seen in the evolution of the MUAC band in Figure 4-19.

4.4.3 Analysis of Measurements

Anthropometric Measurements

We performed two types of analysis for the anthropometric measurements. For the first one, we compare the two automatic measurements with each other, analyzing the absolute value of their difference. This was applied to the three types of measurements. For the second, we compare the automatic measurement with the manual measurement entered by the ASHA, and analyzing the error between the 2, assuming that the manual measurement is the correct value.

For this purposes, box plots were generated to visualize the distributions of the differences. The data for each measurement was filtered from outliers by comparing each data point to the guidelines provided by WHO [25]. When a data point was not within 4 standard deviations of the mean, it was ignored from the analysis. The plots can be found in figs. 4-20 to 4-22. The means, medians, 25th and 75th percentiles were also extracted and are shown in table 4.5.

	Length (cm) |App1 – App2|	Weight (Kg) |App1 – App2|	Weight (Kg) App – ASHA	MUAC (cm) |App1 – App2|	MUAC (cm) App – ASHA
Median	1.08	0.16	0.01	0.83	0.51
Mean	0.5	0.1	0.03	0.24	0.17
25th Percentile	0.23	0.04	-0.08	0.08	-0.12
75th Percentile	1.26	0.2	0.13	0.71	0.48

Table 4.5: Distribution of measurements

While in general we found good agreement within automatic and manual mea-

surements, it is important to note that the widespread distribution of the error is far from perfect and not suitable for a massive deployment yet.

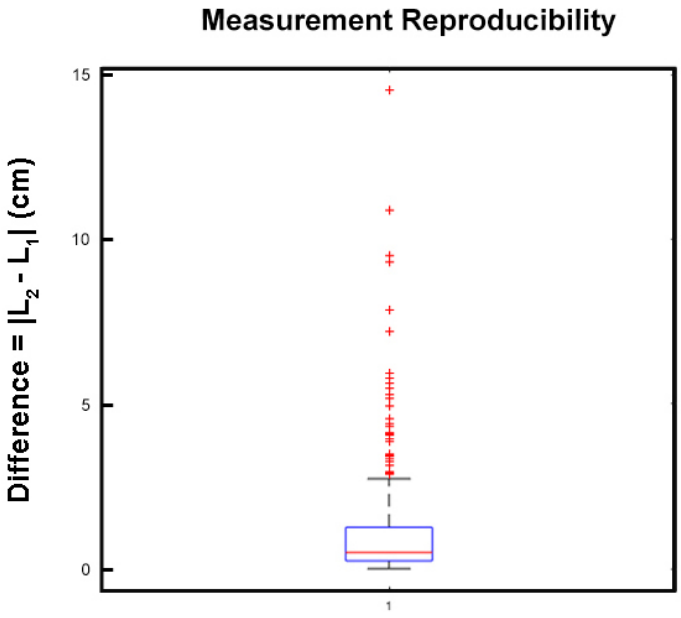

Figure 4-20: Box Plot of length measurement reproducibility, defined as the absolute difference of consecutive app measurements L_1 and L_2

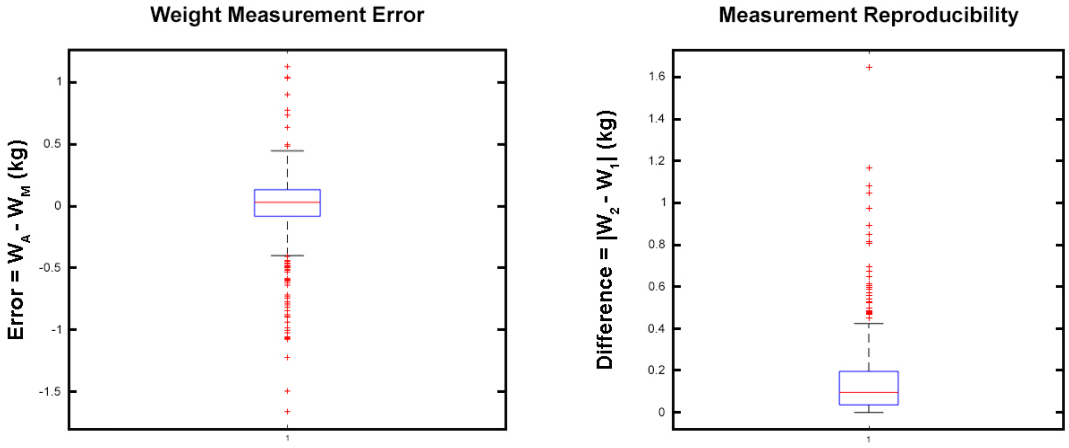

Figure 4-21: (left) Box Plot of Weight Measurement error, defined the difference between the automated measurement, W_A, and the manual measurement, W_M; (right) Box Plot of weight measurement reproducibility, defined as the absolute difference of consecutive app measurements W_1 and W_2

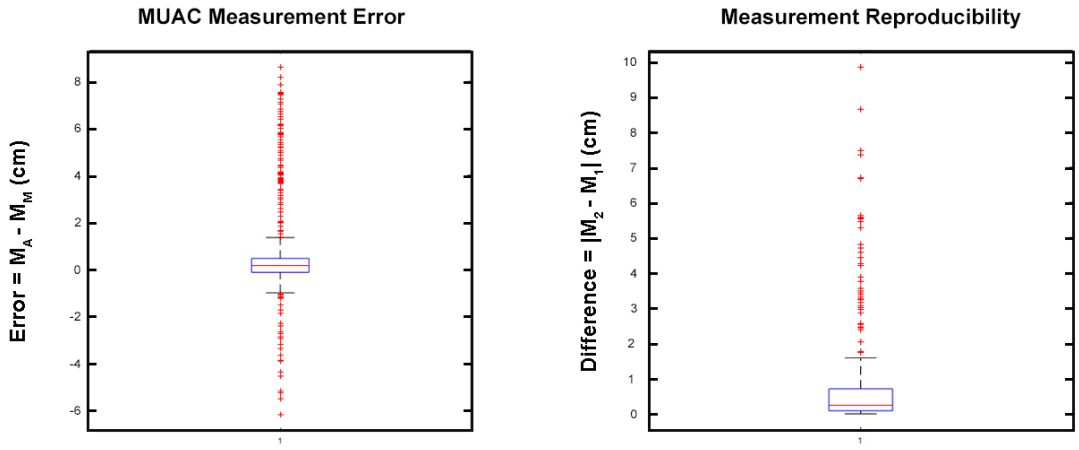

Figure 4-22: (left) Box Plot of MUAC Measurement error, defined the difference between the automated measurement, M_A, and the manual measurement, M_M; (right) Box Plot of weight measurement reproducibility, defined as the absolute difference of consecutive app measurements M_1 and M_2

Timing

We collected the time that it took each ASHA to perform every single measurement. The stopwatch started when the ASHA clicks the measurement button in Baby Naapp, and stopped once the measurement had returned to the summary screen. The expectation for such metric was to verify that once ASHAs became comfortable with the smartphone and the mHealth tools, the effort and time for each measurement would decrease. We were able to verify this hypothesis, by compiling all the timings for each ASHA and fitting an exponential on the data set. Figure 4-23 shows the example of ASHA 8, who has a typical distribution of timings shown as scattered points in blue, while the exponential has been fitted in red. Figures 4-24, 4-25 and 4-26 compile and superimpose the learning curves for each ASHA. The was majority of the curves decay over time, while only few behave in the opposite manner. The MUAC measurement had the least steep slope, while the blanket tool was the one with the most improvement. An interesting observation is that ASHA 3 was the slowest learner for all the tools, while ASHA 6 was always the fastest. One can hypothesize that prior experience with a smartphone might have played an important role in such behavior.

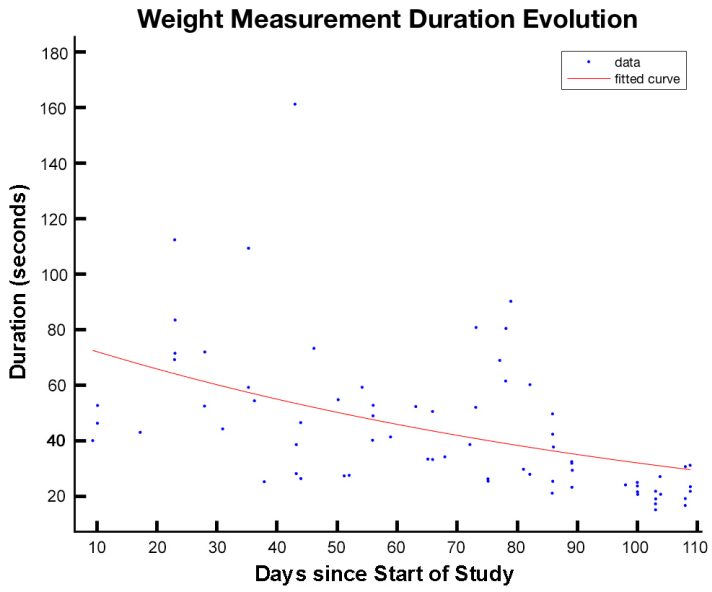

Figure 4-23: Example of Timing Evolution for ASHA worker collecting the weight measurement

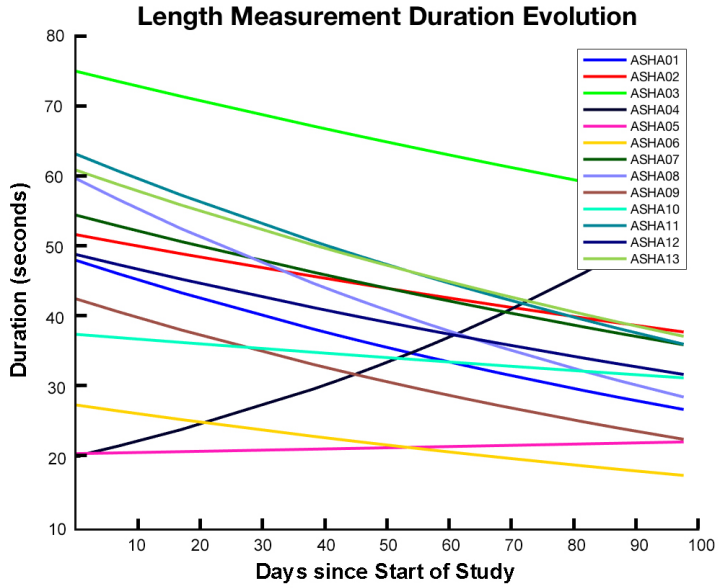

Figure 4-24: Timing Evolution of Length Measurements for all ASHA Workers

Mapping

The development of a mapping interface that displays all the places where measurements happened is in progress in conjunction with undergrad student Jiayu (Jenny) Xue. This project uses the GPS coordinates collected during each measurement to

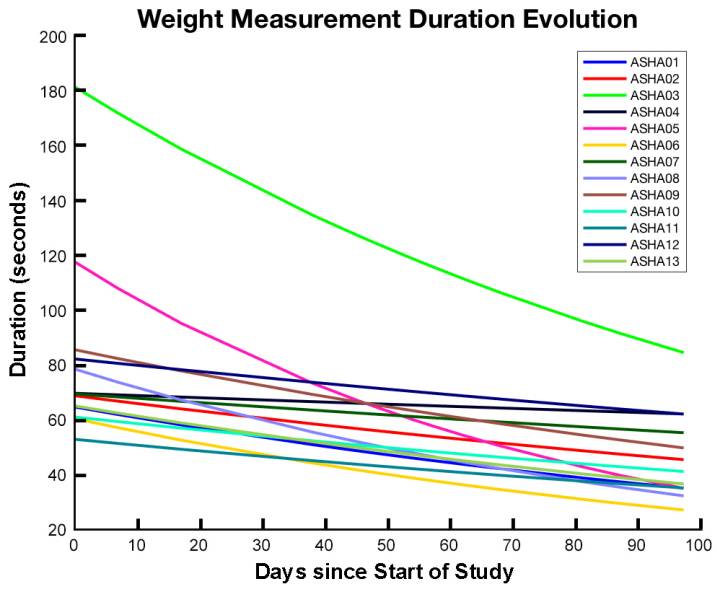

Figure 4-25: Timing Evolution of Weight Measurements for all ASHA Workers

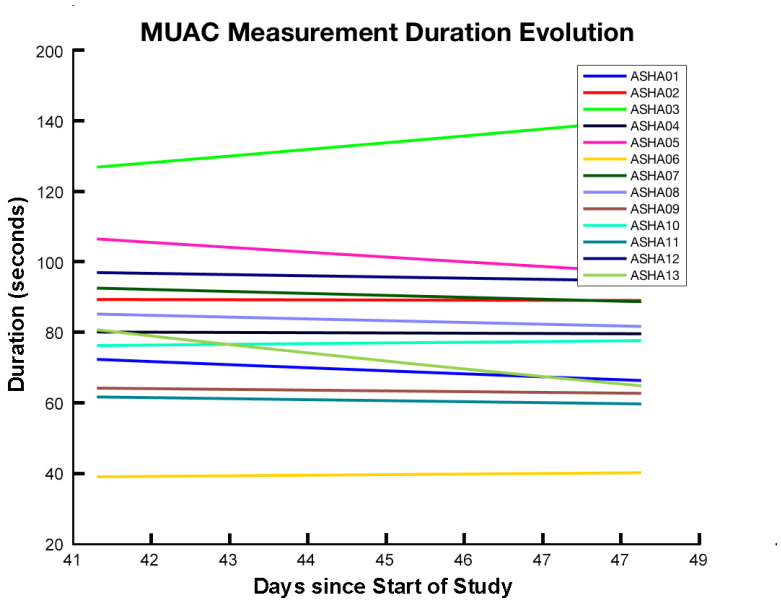

Figure 4-26: Timing Evolution of MUAC Measurements for all ASHA Workers

show the spatial distribution of the areas that ASHAs had to cover. It displays the number of children that were visited in a particular neighborhood, and it indicates whether or not there are malnourished children in that neighborhood. An important takeover from this experiment was that ASHAs learn how to turn off location services and since there was no enforcement or requirement of enabling location services in

Baby Naapp, they inadvertently obscured some of the location information. For cases in which the location was not available, the map tool tries to see if the same baby had another measurement, either before or after, on which the location was available, and if there is, it reuses that location. If there are no other visits, or if there is no location available for any of the visits, the measurement is defaulted to the Delhi Government Dispensary.

The government of India showed great interest in such tools, as one of its possible applications is to target specialized interventions based on the spatial distribution of malnutrition and other afflictions. A preliminary version of the tool can be seen in Figure 4-27.

Figure 4-27: Map displaying the areas on which measurements and visits were made during the study

4.4.4 Feedback From Focus Groups

Two focus groups were conducted after the initial training. One July 8th and the second one on September 9th. The sessions were used to have open conversations with the ASHAs, answer their questions and listen to their suggestions. The meeting

also served occasions for rolling out updates of the apps, and for verifying that the data on the phones was synchronized with the data on the server.

The main topics discussed in the focus sessions can be summarized in the following points:

- They expressed their discomfort with the initial version of the scale app and the difficulties because of the size of the blanket.
- They were worried about using the phone for personal uses, such as taking pictures of their children.
- Some ASHAs shared clever ways in which they had used the smartphones: One used the camera to document a share a record with her superior.
- They shared that many parents were excited with the new technology: "We are feeling good... Until now all the parents who would otherwise take their newborns to a private facility, have become interested and have started allowing us to take their children's measurement. This has increased our credibility. We have gained their trust."
- They expressed their interest in having similar tools for maternal health.
- They were asked about the difficulties and the learning curve of using mHealth tools.
- They discussed how they feel about being an ASHA in general.

4.5 Discussion

4.5.1 Challenges and Limitations

We have already discussed the challenges of dealing with curved surfaces, especially when the curvature of such surfaces is not fixed. Another general limitation lies, not only on the pixel resolution of the camera itself, but also on the different unpredictable

processing steps that phone companies do to "improve" the quality of their screens. While such steps are probably pleasant for the consumer, from the point of view of using the camera to capture a real scene to information, the less processing and the more raw of an image we can obtain, the better.

Vuforia posses some limitations as well. It takes control of the camera resource and camera preview, allowing us to interact with them only via its interface. It is to be explored if in the premium version, the developer is given more liberties over the source code or more faculties to control the camera.

Of course, the resource-constrained nature of the tools bring additional external challenges, starting from dark and small rooms on which our tools may not operate reliably, unforgiving weather conditions that can endanger the integrity of the electronics of the cellphone, deform or stain the blanket or even melt the sticker on the scale.

Another recurrent challenge and a already established pain point is the lack of connectivity in rural areas. If we want to automatically sync and host the data of the phones into a server, a significant amount of work needs to be done to guarantee that in such conditions with limited or intermittent connectivity, data sharing can be achieved reliably without compromising the integrity of the information.

Finally, the large variability of Android phones in the market prevents us from guaranteeing that our apps will perform flawlessly in their system.

4.5.2 Future Improvements

Each of the anthropometric tools needs to be improved before they can be massively deployed into the developing world. If the solution to the problems that were discussed in the past two chapters were known by the author of this work, they would already be implemented and presented. However, there are some ideas of refinements that can be made and they shall be discussed.

For the Blanket app, additional pieces of hardware need to be added to hold the feet and the head of the baby. Such pieces could have additional targets to lock an accurate reading of the length. The software can also be improved. The simple blob

detection that it performs now is not enough to handle non ideal scenes. A better algorithm can be develop to understand the shape and the color of the baby.

The scale has seen a great improvement by transitioning to the cylindrical target. However, more work can be done. The unpredictable auto white-balance and dynamic range algorithms that are performed by the phones sometimes confuse the tracking algorithm that is trying to differentiate between light and dark. Again, the use of an additional target would be ideal in this scenario. Estimating the distance between two targets seems more reliable than differentiating between black and white in an environment on which such concepts are relative.

Also the MUAC needs to have reinforcement in its algorithm. The measurement is vulnerable to noise by the slightest rotation of the band or the phone, and even by the shadows than can be naturally casted on it. The phase based algorithm is a good base to create an educated guess, but a viable options to solve this problem is to reinforce it: OCR, Color recognition, spatial gradient analysis are just some examples of techniques that could improve the MUAC algorithm.

Finally, to move forward there is a deep need into gaining control over the resources of the phone. Whether it is the camera itself, or the freedom of managing memory, the addition of more complex algorithms will require the developer to have the ability to write code natively.

Figure 4-28: All the ASHAs and researchers captured right after the first focus group

Chapter 5

Using Machine Vision for Physiological Measurements

5.1 Motivation

According to the Merriam-Webster dictionary, physiology is "a branch of biology that deals with the functions and activities of life or of living matter (such as organs, tissues, or cells) and of the physical and chemical phenomena involved" [26]. Hence, is not possible to perform a complete health assessment without collecting and understanding physiological measurements. As a matter of fact, an article published in the Science Translational Medicine Magazine [27] established a causal relationship between early physiological responses to predict later illness severity and even mortality in infants.

Examples of physiological measurements include heart rate, respiration rate, blood sugar levels, red blood cell count, oxygen saturation, etc. What these measurements have in common is that the gold standard for their assessment involves expensive laboratory analysis machines, and invasive or direct contact methods of collection. Arguably, the holy grail of physiological measurements in the developing world lies in developing inexpensive, contactless sensors that can extract the same information with minimal quality degradation. In this context, a natural expansion from our work measuring anthropometry with the camera is to develop more advanced tools that

can understand subtle changes in the shape, the color, or the movements of a person in the camera frame to extract information about its physiology.

5.2 Some Examples of Prior Work

A lot of work has been done to explore the area of contactless physiological measurements. A method called Eulerian Video Magnification (EVM) [28], was developed here at CSAIL. It makes use of small temporal changes in the frames of a video that are imperceptible to the naked eye, to extract, amplify and reveal information hidden in such variations. With this method, researchers were able to display blood flow into the face as well as amplified range of motion during the respiration of a baby, and measure heart and respiration rate. The method has been used as the underlying process for implementations in different devices such as the Microsoft Kinnect [29], to explore and advance the area of contactless heart rate detection. Unfortunately, as with video plethysmography, the technique is fundamentally limited to a mounted fixed camera.

Color analysis and colorimetry are also important techniques used to measure physiology. In section 2.1, we showed a concrete example [9] of this technique, that compares different tonalities of red to estimate hemoglobin levels and detect anemia. Color analysis can also be used to estimate the heart rate and heart rate variability: Our group developed a simple tool that is able to capture a PPG signal by using the flash as the light source and the camera as the optical sensor [23]. The Ubiquitous Computing Lab (ubicomplab) at the University of Washington has also done extensive research on this method, using it to extract hemoglobin levels in [30] and [31]; and combining it with the accelerometer to extract blood pressure [32]. ubicomplab has also explored machine vision techniques to screen for jaundice ([33] and [34]), brain injury [35], and intraocular pressure [36].

The work that certainly served not only as inspiration, but also as a code base for the thermal breath analysis work exposed in section 5.4, was the implementation of a smartphone video plethysmography method described in [8]. While for that particular

project physiological measurements were unfeasible for non-static persons because of the low signal-to-noise ratio on moving scenes, in the case of respiration rate thermal analysis, lower frequencies allow for the scene to be accommodated in order to always obtain strong breathing signals. Thermal imaging and its applications are explored in the following sections.

5.3 Applications of Thermal Imaging

So far, we have only explored the uses of machine vision in the visible light spectrum, but the infrared (IR) spectrum contains information that can be as rich and at times invisible to the naked eye. Since most of the thermal radiation emitted by objects near room temperature is infrared, IR imaging provides us with an array of temperatures that can be thought of as a different color representation of an image that contains vast amounts of new information. For these reason, we decided to expand our work into the IR region.

5.3.1 Screening for infections in babies

In a newborn baby, many regulatory systems begin to function outside of the mother's womb. One of these is the baby's own mechanism for regulating its temperature. The thermal regulation of a neonate is critical for survival and its measurement can be an indicator of pathophysiology [37]. This is quite difficult for a baby, since it has a small mass and small amount of blood. Studying the thermal regulation of a newborn provides important information regarding the health of a baby as well as it's maturity. For example, measuring the temperature of the baby's limbs compared to its core temperature is an indication of how well a child is able to regulate its temperature. Thermal analysis is also a useful indicator for infection, circulatory shock and sepsis [38], or pneumonia [39] — with each of these having its own signature.

5.3.2 Screening for pulmonary diseases

A study published in the BMJ Open Journal showed a proof of concept that explored the potential of thermal imaging as a valid, innovative and inexpensive technology for diagnosing bacterial pneumonia [40]. The study analyzed the similarity of chest X-ray and thermal imaging on patients that presented focal pneumonia. It concluded that thermal imaging can detect pneumonia with high sensitivity and negative predictive value in a consistent way with the X-Ray method.

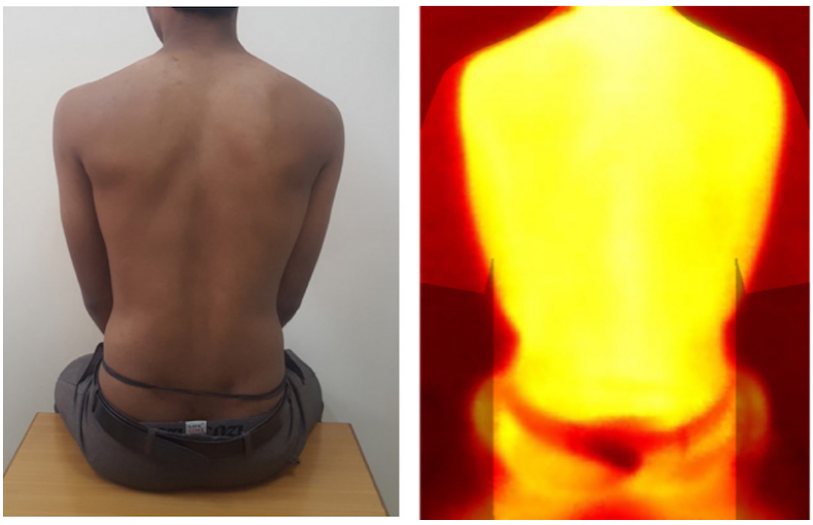

Figure 5-1: Example pictures collected with the thermal tool. The thermal image has a superimposed semitransparent torso silhouette to demonstrate how the image looks when it is being captured

With such capabilities, our group is also interested on developing a thermal screening tool that can detect a localized infection, such as tuberculosis. For this purpose, our group has partnered with the Chest Research Foundation (CRF) in India and we have developed a data collection tool that uses a USB thermal module that plugs into the smartphone (For details on the thermal module, refer to 5.4.1). The app collects both, the visible and the thermal image (as seen in Figure 5-1), and provides the clinician with a workflow to collect images of the patient from 8 different positions: posterior, left side, anterior and right side, all of them with an exhale and inhale variation. To guide the clinician's process on capturing a thermal images, different semitransparent silhouettes (Figure 5-2) are superimposed on the thermal camera

preview and instructions are displayed in each step of the process.

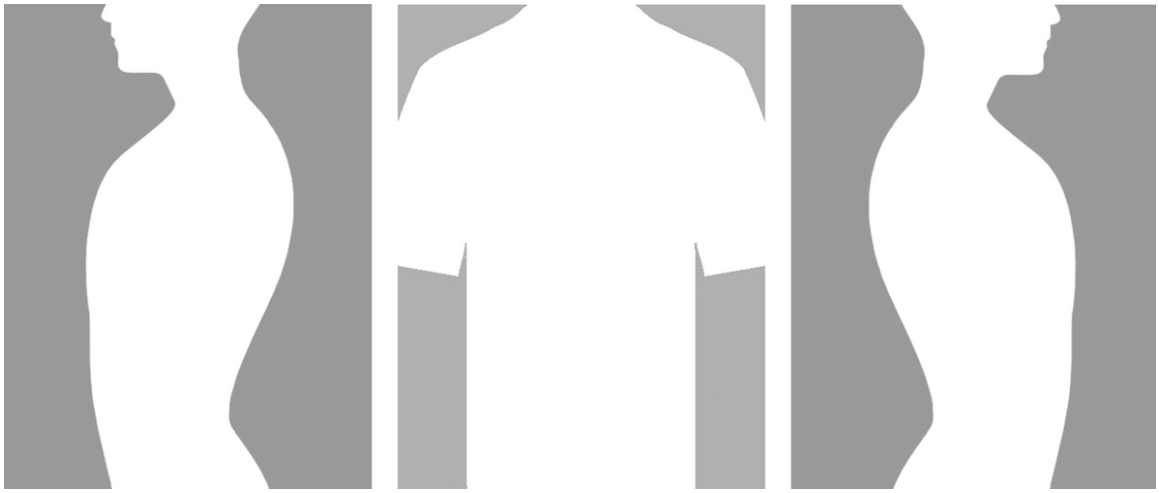

Figure 5-2: Images of the silhouette filters used to guide the clinician's process on capturing a thermal image

5.4 Use of Thermal Imaging For Breath Monitoring: A Lab Study

Measurement of human breathing is a fundamental component of clinical diagnosis and disease screening. Breath rate and breath rate variability are important parameters that can be used to help diagnose and monitor a variety of disease conditions including pneumonia and chronic obstructive pulmonary disease, and can also be used as part of mental health assessment as an indicator of stress or relaxation. The ability to perform non-contact assessment of breathing is particularly useful for health screening of infants and neonates, very old frail patients, or patients with infectious diseases, where it is not practical or advisable to use any device that attaches to the patient.

We are developing a mobile app that uses thermal imaging and machine vision to automatically measure respiration rate (RR) and respiration rate variability (RRV). The airflow generates measurable fluctuations of temperature in the nostrils and in between the nose and the mouth. While other groups have demonstrated the use of

thermal imaging to measure the respiration [41], our approach utilizes low cost, low resolution camera to account for resource-constrained environments, and explores a variety of configurations to maximize the signal strength.

We have partnered with the Swamy Vivekananda Yoga Anusandhana Samsthana University (S-VYASA) in Bangalore, India, to explore the possibilities and limitations of the technology applied to contexts of global health and, in particular, yoga. A preliminary study was conducted and its details will be discussed in the following sections.

5.4.1 The Thermal Module

For implementation, we made use of a commercial USB thermal imaging module designed for use with mobile phones (Seek Thermal Compact), which has a pixel resolution of 206x156 and an advertised frame rate of about 9 Hz (Figure 5-3).

Figure 5-3: Thermal camera module used with thermal apps

To understand the actual behavior of the camera we generated plot 5-4, which contains the horizontal axis the elapsed time, and on the vertical axis the number of frames that have become available. They are divided by the average frame rate to illustrate how, for a constant sampling frequency, the plot would follow the line $y = x$ in red, but since the sampling is not constant and it stalls periodically, the actual plot, in blue, deviates from that of a constant sampling frequency.

For development we use the development SDK Seekware, which was provided by the manufacturer of the thermal module. The SDK provides sufficient functionality such as absolute temperature arrays for every frame and a variety of color mapping from infrared to create images in pseudo colors.

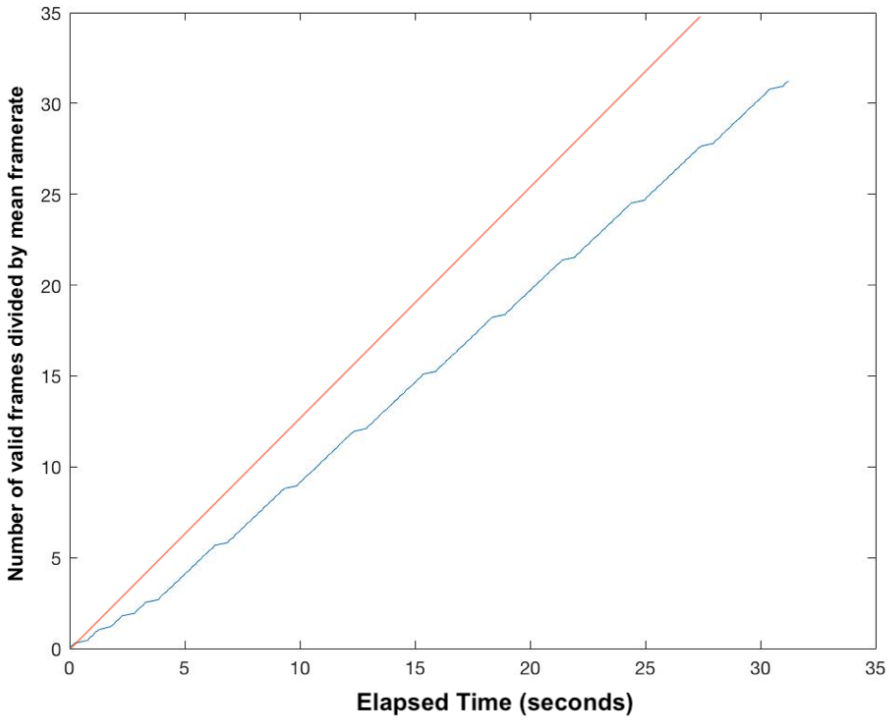

Figure 5-4: Frames Available vs. Time plot. In red the $y = x$ plot, and in blue the frames signal.

5.4.2 Thermal Breath Analyzer Mobile App

The prototype app was developed for Android 7.0. It identifies a ROI containing the nostrils, and samples the absolute temperatures of a Gaussian-weighted neighborhood around each individual nostril. The sums of those neighborhoods are passed into the signal processing pipeline described in section 5.4.3. The filtered signals are displayed on the top of the screen (as seen on Figure 5-5) and the respiration and the respiration rate may be extracted, by means of Fourier analysis or peak finding.

The goal is to evolve this prototype into a variety of tools. One can result into a coach that can assist users into proper breathing and relaxation. Another one is the use of the tools as health assessment means to capture RR and RRV.

5.4.3 Signal Processing Pipeline

The signal processing pipeline of the app can be summarized in the following steps:

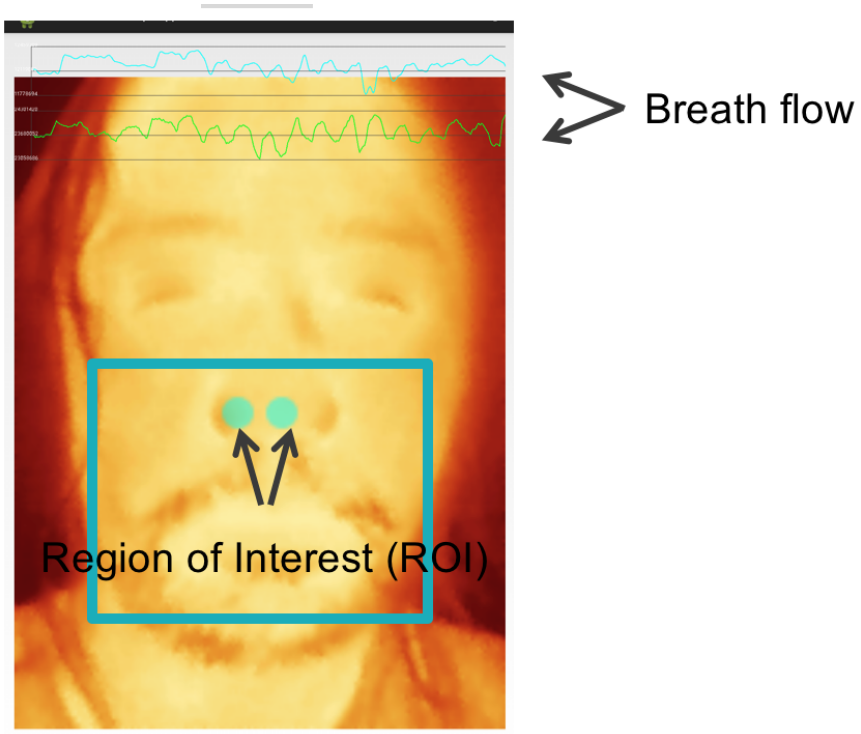

Figure 5-5: Layout of the Thermal Breath Analyzer App

1. Detection of ROI

2. Extraction of temperature neighborhood values (in float which is the data type provided by the Seekware SDK)

3. Conversion to integer type arrays (for efficiency reasons)

4. Calculation of the sums of temperature values over the left and right nostrils neighborhoods

5. Time Series Generation that includes resampling to an empirically chosen mean frame rate

6. Application of FIR Low-pass filter with a cutoff frequency higher than normal respiration rates

7. Time Derivate Calculation using the 5 point rule

Since the temperatures of the ROIs fluctuate based on inhalation and exhalation, we predict that the time derivative of the temperature signals can approximate air flow rates and calculate timing parameters.

5.4.4 Data Collection and Preliminary Analysis

Figure 5-6: Pictures of the positions that were explored for the Thermal Breath Analysis app

	Breath Rate using Respiration Belt				Thermal Camera [Book Position]			
Subject	Normal Breathing	Deep Breathing	Fast Breathing	Shallow Breathing	Normal Breathing	Deep Breathing	Fast Breathing	Shallow Breathing
1	10.44	8.79	59.33	24.94	9.38	3.22	28.29	27.55
2	6.78	15.94	13.95	30.95	5.86	2.93	15.83	32.83
3	11.16	5.67	23.59	10.33	10.55	2.63	24.62	10.26
4	11.3	6.21	32.48	33.39	11.13	5.86	29.90	32.24
5	19.85	24.41	31.14	37.01	18.46	22.27	29.31	40.16
6	14.16	8.89	47.05	27.16	15.53	2.63	49.54	26.09
7	20.82	7.06	30.16	29.25	20.81	3.81	31.07	27.26
8	23.47	8.21	37.32	30.04	23.45	4.98	2.64	35.47
9	17.23	10.76	29.56	27.98	16.70	11.13	28.43	27.85
10	11.3	4.39	33.7	26.44	10.55	4.39	35.18	25.50
11	9.67	8.04	24.86	24.36	9.38	3.22	27.85	27.85

Table 5.1: Respiration Rate Estimates

Our device was tested with eleven human subjects in a clinic and validated against a gold-standard piezo-electric respiration chest belt, that responds linearly to changes

in length, and for three intensities of breathing (normal breathing, shallow breathing, and deep breathing) as well as two different position of the head: camera straight ahead or "mirror position", camera slightly below or "book position" (20-degree angle). The mirror position simulates the position of a user looking themselves at the mirror. It generates a noisier signal than that of the book position, but the configuration feels more natural for the user. On the other hand, the book position simulates the position of the user reading a book. It generates a stronger signal since the camera is looking at the nostrils from a more direct angle. The shortcoming of this position is that ideally, the user needs to keep his head straight and resist the temptation of bending his neck to look at the thermal image. Pictures of both configurations can be found in Figure 5-6.

With the thermal camera placed in the book position, the results from the mobile phone tool had good agreement with the clinical device, with preliminary trivial Fourier analysis resulting in mean absolute errors of 0.64, 3.81, 7.46, 1.99 bpm, for normal, deep, fast, and shallow breathing, respectively. The estimated rates can be found on table 5.1.

An example of different representations of the same temperature signal is presented in the plots of Figure 5-7. For time series of normal breathing, it is possible to obtain clean signals with well defined peaks, whereas for the Fourier domain, the dominant frequency is usually that of the respiration rate.

5.4.5 Face Tracking and Detection of Region of Interest

Naturally, the book position was chosen to develop the automatic detection of the region of interest. For the initial detection of the face, we incorporated OpenCV into our mobile app and used their pipeline [42] to train a HAAR Cascade that can recognize faces in the infrared spectrum, from the book position orientation. Examples of the positive images used for training can be seen in Figure 5-8. For detection, we used the Viola—Jones [43] and [44] object detection framework that is also built-in the OpenCV Android SDK.

Since the resolution of thermal camera is not too high, we were able to complete

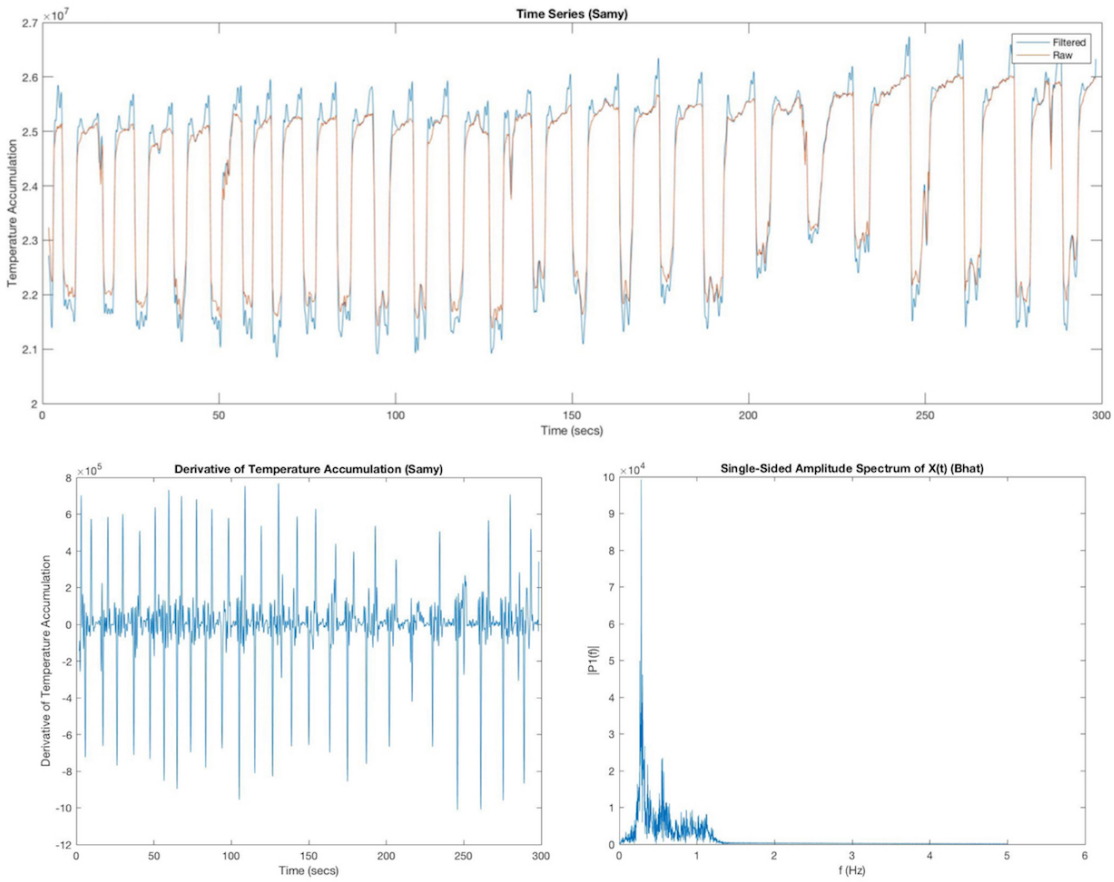

Figure 5-7: Plots of different of temperature signals. On the top, the raw and filtered temperature signals. On the bottom left, the time derivative of the top plot. On the bottom right, the Fourier representation of a temperature signal

8 stage training with 200 positive images and 610 negative images. The face tracking works reliably for static faces, but it can be improved by adding some persistence of the tracking based on the assumption of slow movements.

It is important to note that the OpenCV tracking takes place in the color mapped image and not on the thermal array. We have chosen the particular color mapping that is shown in the example training photos, because it provides a very good representation of the features that were are interested on. In other words, for this color mapping named *SeekwareLut.LUT_GLOW*, the nose, the mouth, and the eyes are very distinct in the face.

The rectangular ROI shown in Figure 5-5 was found by the thermal face detector and provides us with an even smaller image to work with. For such small images,

we have decided to abandon tracking with OpenCV, and perform analysis using the temperature array itself.

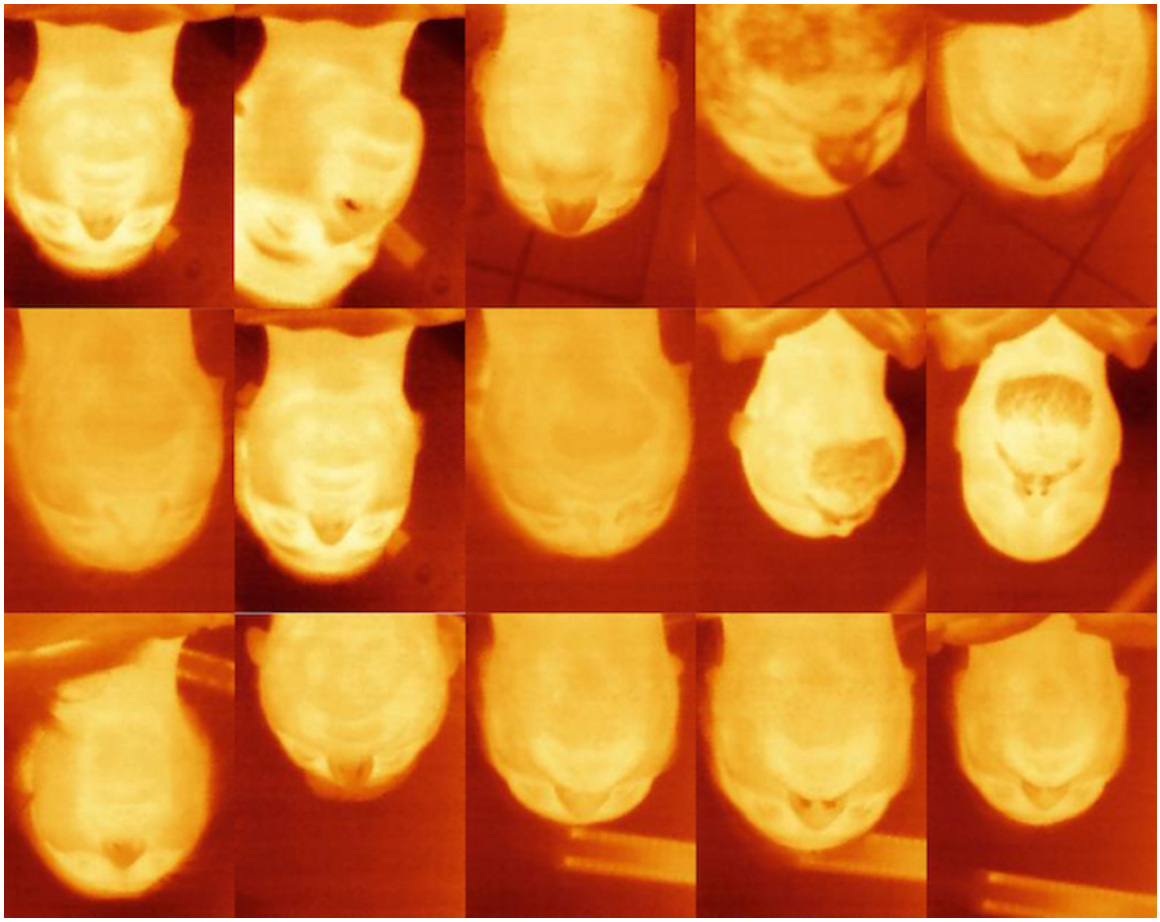

Figure 5-8: Examples of the positive images used for training the HAAR Cascade

5.4.6 Ongoing and future work

Finding the nostrils in the image is the final step that needs to be implemented in order to have a fully connected workflow.

We explored training a HAAR cascade to identify noses in the thermal face, but the method was not very effective. We have tried using Hough Circle Transforms and contour detection, but again the fluctuating signal of the temperature prevents the found areas to persist and are very sensitive to noise.

Moving forward, we will decided to dive into the array of temperature values itself. We will explore using time derivatives and radial derivatives of such array to find the

exact location of the nostrils. The results of our findings will be presented in a future publication.

Chapter 6

Design of a Framework for Scalable mHealth Tools

6.1 Scaling and Deployment Challenges

When a single app needs to be created as an assignment or for the purposes of a specific study, there is no need to give much thought on design decisions for a maintainable and scalable future. As a young group that started developing mHealth apps, our main focus was to deliver products were capable of performing their own specific task and with a very limited vision of the other apps in its environment. However, as our lab grows and our technology advances, we have increasingly dealt with the challenges of lack of standards, redundant information inputs, and data formats that are convenient for specific analysis, but not easy to share.

A design on which each app is focused on a specific task is attractive for the developing world, since it comes free from unnecessary and often unwanted features. However, such design does not contradict a scenario in which a specialized app can live in the same environment and share resources with similar ones.

6.2 A general App Framework

Even as the verticals of our lab became clearly defined (the 3 largest ones being Pulmonary, Cardiovascular and Neonatal and Child Health), it also became clear that a new framework had to be designed that could:

- Abstract and Modularize common components from the different apps.
- Create rules that would allow apps to communicate within each other and share mutually needed information.
- Pave the way for integrating with technologies that are becoming increasingly ubiquitous, such as cloud based computations and Internet of Things.
- Preserve the independence of apps specialized in particular measurements.
- Update to and follow UI and code standards that are the Android Developer Guide.

The first step was to create a home directory on which all of the apps of our apps would live. The name of our home is MobileTechLab, and in this way the paths of all of our apps were change from /SpecializedApp to /MobileTechLab/SpecializedApp.

6.2.1 Abstracting the common the denominator

A summary or data aggregator app was the first attempt of our group to move into being able to share and combine the measurements of different specialized apps. It was presented as a pulmonary screener tool in [45]. The app, whose main workflow was a questionnaire, gave the user the option of plugging into a specialized app that used the phone camera to automatically record the reading from a peak flow meter (PFM) device, instead of forcing the user to enter it manually. The interaction is simple: the pulmonary screener creates an intent to call the PFM activity, and populates the intent with the information that the PFM requires to analyze its measurement. Once the PFM measurement is completed, the activity returns to the Pulmonary Screener with the values calculated by PFM activity. It is important to note, that for

this workflow to work the Pulmonary Screener asks the Android Activity Manager for someone to perform the specialized task. The only way in which the Android Activity Manager can know about such app, is if the app explicitly declares in its manifest a listener for such request.

Baby Naapp and all the NCH tools follow this model, but in order to make it scalable, the way in which shared data is dealt with needs to be standardized. For example, the PFM app requires the height of the patient, but if given the ID of the patient, it could potentially handle the case of a missing height. Does it make sense for the PFM app to crash its operation if a height value is not passed? Should Pulmonary Screener be responsible of sending the correct and complete data to PFM or is it a better alternative to allow the PFM to retrieve the height by itself? The answer to the last question is that it should support both options.

In terms of efficiency, if the Pulmonary Screener needs to open the patient profile to retrieve metrics from the patient, it makes sense for it to keep the most commonly needed information readily available. For each specialized measurement app, there should be a contract specifying the pieces of information that it requires when called externally. If the caller app is not able to pass one of the requirements, the specialized app should be ready to explicitly ask the user to enter it, or try to retrieve it from the shared profile information. But before such behavior can be achieved, first the apps need to maintain shared profile information.

6.2.2 A shared Profiles Directory

Of course, different apps require different specific information, but most of them usually share common fields such as IDs, dates of births, and genders. In the past, each app had maintained its own profile in a format that could only be understood by that particular app, but we are leaving behind such practice. A Profile Class was created to standardize the way in which profile information is stored, while the Profile class gives autonomy to the developer to incorporate fields at will, it comes with a prior knowledge of fields that are commonly used. Profiles are stored and retrieve in the form JSON files. The handling of JSON is done via the library GSON, which

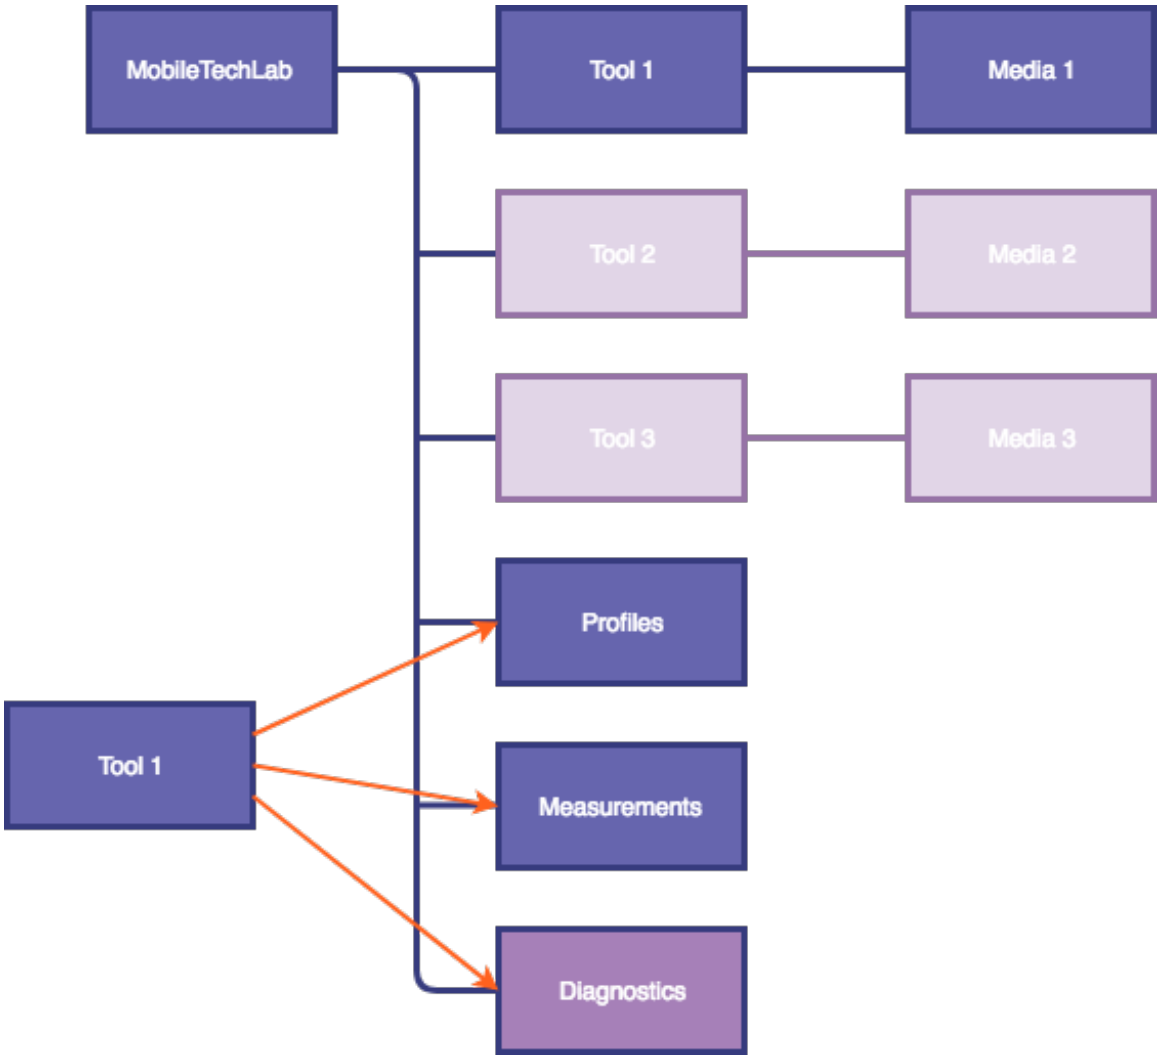

Figure 6-1: Diagram of the new Folder Structure: It shows how each tool has access to the share profiles, shared measurements, and shared diagnostics (new), while at the same time it separating the information of its own media files, such as pictures of ROIs, or internal calibration numbers

provides a handy API to convert the file into a Java key—value data structure. All the profiles are saved with the name of their unique ID and in a location that is known to all the apps: /MobileTechLab/Profiles.

Dealing with missing data now becomes very methodical. If a required field is missing for a specialized app, it can try to retrieve the value from the profile. If the field is not found in the profile, the app should explicitly ask for it. If the field is a fixed value, like the height of an adult, it can now be added to the profile to be available for any other app that requires it. If it is not fixed, like the height of a child,

the field may be used accordingly and disposed.

6.2.3 A shared Measurement Directory

The scenario was similar for measurements of each individual app, with them being isolated from each other. Here as well, the first natural change was to migrate the format JSON. Useful info such as units, timings and location were included in each measurement, but now all types measurements have started concurring into the same file, it is important to establish standards and conventions. Each type of measurement is being assigned a unique integer identifier, which is included in the measurement JSON Element. We also identified that another concept higher in hierarchy than that of a "measurement" is often used: this is the concept of a "visit". Measurements can be grouped in visits, and while more than one visit can happen on the same day, it is often assumed that a different date implies a different visit. In this way, a single measurement file is created for every user in the MobileTechLab ecosystem. The file is saved with the name of the unique ID of the patient in the directory /MobileTechlab/Measurements. As a result, measurement files are composed of a list of JSON objects that represent visits, and each visit contains a field with a list of the measurements for that day. An example of a section from a sample measurement file can be seen on Figure 6-2.

An additional benefit of transitioning to JSON-formatted measurements is to automatically become more cloud and server friendly, as JSON is the standard used for most modern web APIs.

6.3 Case Studies and Examples

6.3.1 Neonatal and Child (NC) Screener

While Baby Naapp was an app specifically designed to meet the ends of the CHW Study, the Neonatal and Child Health Screener App (NC-Screener) is a spin-off of Baby Naapp that adopts all of the design paradigms discussed in the chapter and can

Figure 6-2: A sample section of a measurement file

be thought as the product version of Baby Naapp. All of the individual anthropometric measurements were also updated to follow the new profiles and measurements standards. However, each of them still has its individual directory in MobileTechLab. The anthropometric apps, maintain a Media directory on which the ROIs may be saved for comparison with the measurements extracted. In general, each apps folder may save diagnostics or information that is specific to that app and does no need to be shared.

Additionally, we added an extra results screen that plots the anthropometric measurements in graphs that contain growth charts provided by WHO. While CHWs usually use printed copies of such charts to plot an individual measurement data point, our version allows for digital longitudinal tracking of the data, an important metric to assess for malnutrition.

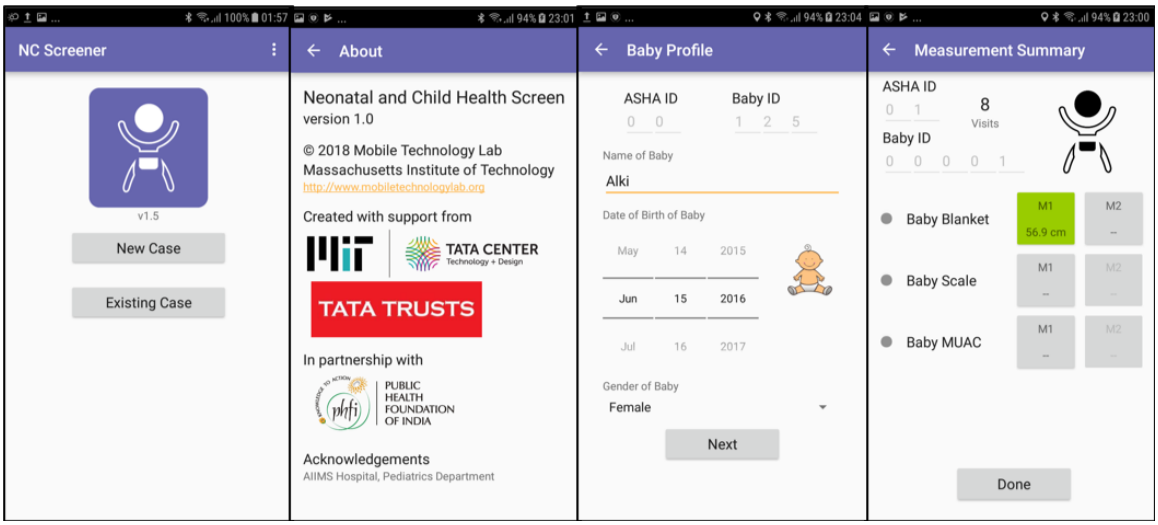

Figure 6-3: Screenshots from NC-Screener

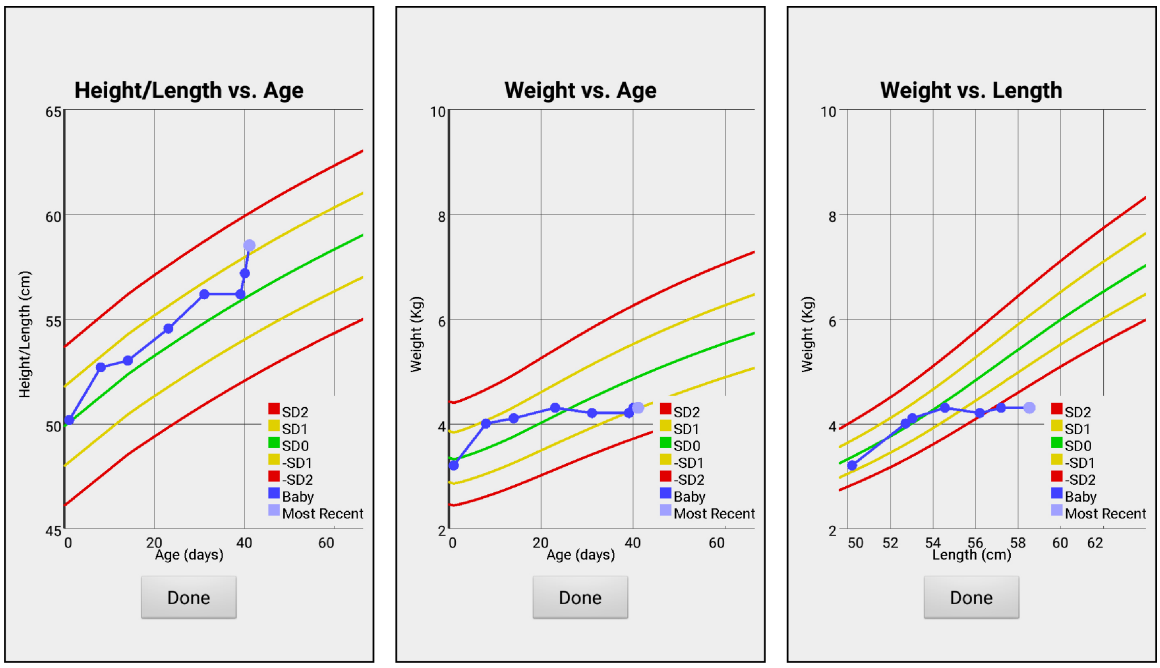

Figure 6-4: New Plot Screens for Anthropometric Measurements

An immediate consequence of the new format of the measurement file is the ability to easily combine measurements from two different tools as observed in the Weight vs Length Plot.

6.3.2 Pulmonary Screener

A Pulmonary Screener v2.0 was developed as well. The pulmonary apps have already adopted the use of the common profile directory and paradigm. Since Pulmonary Screener v2.0 is mainly a questionnaire, it was the perfect scenario to ask the profile for answers to some of questions such date of birth or gender. It is possible that time series results stay in the individual folder of the specialized apps, while metrics extracted from them might end up in the measurement file.

An additional folder that is planned to be introduced is a Diagnostics directory, which would contain files similar to the measurement ones, but diagnostics would combine the information from different measurements to categorize the patient, diagnose an affliction, or give a recommendation, that would get stored in the diagnostics file. While NC-Screener does not generate any diagnostic data yet, pulmonary screener does, and it will continue to adapt to this newer protocols.

The fate of the roll-out of this framework lies on its near-future proper documentation.

Chapter 7

Conclusions

7.1 Next Generation of mHealth Tools

With the release of ARKit by Apple and ARCore by Google over the past year, and with the continued growth of older AR Development Kits such as Vuforia, it is clear that the smartphone industry is moving forward into a extensive exploration of such technologies. In addition, the inclusion of dual cameras, infrared sensors, and powerful processors that focus more and more on image processing, only confirm the clear direction in which mobile hardware manufacturers are heading. With such power, it is tempting to assume that computations will only keep becoming more complex in the smartphones.

However, pushing from the other direction, there is cloud computing which in principle has an unbounded number of resources for computation. Its limitations lie on the transfer speed between the phone and the servers and on the sometimes unreliable Internet access. However, as future technologies such as 5G become more of a reality, the limitation of transfer times steadily fades away.

Future technologies must meet halfway and create a seamless interaction in which users do not realize that computations may not happening locally. The challenge for developers will be to identify where the new bottlenecks are formed.

For users in remote areas, there will also be a shift in the way in which they adapt to new technologies. Connectivity is part of a positive feedback loop system

that catalyzes the adoption of technologies, and the new technologies promote more connectivity.

On a different point, wearable sensors keep becoming more inexpensive and popular in the developed world, but they will reach a saturation point in which their new capabilities will only be incremental updates. At such point, wearables will need to either go in, and become one with the user as implants or more attractively, go out and become a natural part of the ambient. It is difficult to predict where the next disruption might come from, but with the boom of the Internet of Things, it is safe to assume that it will include an environment or mesh of health sensors that will be present everywhere, from your wifi router detecting if you have tripped and fallen [46], to your smartphone telling you to cough again because it heard a suspicious wheeze.

7.2 Reflection

There are branches of Computer Science and Engineering that can remain buried under layers and layers of abstraction. For Machine Vision, this is not the case and less so for an application as the one explored in this work, with such potential for global impact. And while the math and algorithms behind Machine Vision make it a beautiful field by itself, the countless applications to the real world add a whole new level of satisfaction that I doubt can be met by the purely theoretical or the purely applied.

In this thesis, we have designed and demonstrated a low cost child health assessment kit. We have learned from each iteration of its design and improved it such a way that its operation becomes simpler and its capabilities scalable. We have explored and established the basis for powerful physiological measurement tools, which when ready, will be easily incorporated into the child health assessment. This work has shown that even with "simple" algorithms, there is potential for a significant improvement in the data collection methodologies for CHWs. The word "simple" is in quotes here because our algorithms are only simple due to of the fantastic work done by Machine Vision groups to provide the world with powerful tools that can be

seamlessly connected into larger and more powerful systems. With this in mind, it is only natural to assume, that the methods and tools that we have generated will come to be used as part of a yet larger and more powerful system.

Figure 7-1: Honey Bajaj and Xavier Soriano observing a baby during a visit with and ASHA

Bibliography

[1] World Health Organization and International Bank for Reconstruction and Development and The World Bank. *Tracking Universal Health Coverage: 2017 Global Monitoring Report*, 2017.

[2] R. Fletcher, X. Soriano Díaz, H. Bajaj, and S. Ghosh-Jerath. Development of smart phone-based child health screening tools for community health workers. In *2017 IEEE Global Humanitarian Technology Conference (GHTC)*, pages 1–9, Oct 2017.

[3] Ministry of Health and Family Welfare. Govt. of India. Operational guidelines on home based newborn care (HBNC). (Accessed on 05/25/2018).

[4] National Health Mission of India. List of drugs being provided in asha kit - government of india. (Accessed on 05/25/2018).

[5] The Hindu. 35,000 asha workers to get mobile phones to promote health schemes - the hindu. (Accessed on 05/25/2018).

[6] Errol Ozdalga, Ark Ozdalga, and Neera Ahuja. The smartphone in medicine: A review of current and potential use among physicians and students. *J Med Internet Res*, 14(5):e128, Sep 2012. v14i5e128[PII].

[7] Sana. `http://dev.sanamobile.org/`. (Accessed on 05/25/2018).

[8] R. R. Fletcher, D. Chamberlain, N. Paggi, and X. Deng. Implementation of smart phone video plethysmography and dependence on lighting parameters. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 3747–3750, Aug 2015.

[9] Richard Fletcher, Niccolo Pignatelli, Adrian Jimenez-Galindo, and Suparna Ghosh-Jerath. Development of smart phone tools for printed diagnostics: Challenges and solutions. In *2016 IEEE Global Humanitarian Technology Conference (GHTC)*, pages 701–708, Oct 2016.

[10] UNICEF. Malnutrition rates remain alarming: stunting is declining too slowly while wasting still impacts the lives of far too many young children. `https://data.unicef.org/topic/nutrition/malnutrition/#`. (Accessed on 05/17/2018).

[11] World Health Organization and United Nations Children's Fund. WHO child growth standards and the identification of severe acute malnutrition in infants and children. http://www.who.int/nutrition/publications/severemalnutrition/9789241598163/en/. (Accessed on 05/17/2018).

[12] David G Lowe. Distinctive image features from scale-invariant keypoints. 2004.

[13] Matthew Brown and David G Lowe. Automatic panoramic image stitching using invariant features. 2007.

[14] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, June 1981.

[15] S. Laotrakunchai, A. Wongkaew, and K. Patanukhom. Measurement of size and distance of objects using mobile devices. In *2013 International Conference on Signal-Image Technology Internet-Based Systems*, pages 156–161, Dec 2013.

[16] Dianyuan Han and Chengduan Wang. Tree height measurement based on image processing embedded in smart mobile phone. In *2011 International Conference on Multimedia Technology*, pages 3293–3296, July 2011.

[17] OpenCV. Eroding and dilating — OpenCV 2.4.13.6 documentation. https://docs.opencv.org/2.4/doc/tutorials/imgproc/erosion_dilatation/erosion_dilatation.html. (Accessed on 05/19/2018).

[18] Sylvain Paris, Pierre Kornprobst, Jack Tumblin, and Frédo Durand. Bilateral filtering: Theory and applications. *Foundations and Trends®in Computer Graphics and Vision*, 4(1):1–73, 2009.

[19] Sotirios Fouzas, Kostas N. Priftis, and Michael B. Anthracopoulos. Pulse oximetry in pediatric practice. *Pediatrics*, 128(4):740–752, 2011.

[20] T. Duke, R. Subhi, D. Peel, and B. Frey. Pulse oximetry: technology to reduce child mortality in developing countries. *Annals of Tropical Paediatrics*, 29(3):165–175, 2009. PMID: 19689857.

[21] Qu-ming Zhao, Xiao-jing Ma, Xiao-ling Ge, Fang Liu, Wei-li Yan, Lin Wu, Ming Ye, Xue-cun Liang, Jing Zhang, Yan Gao, Bing Jia, and Guo-ying Huang. Pulse oximetry with clinical assessment to screen for congenital heart disease in neonates in china: a prospective study. *The Lancet*, 384(9945):747–754, Aug 2014.

[22] Martin Weber, John B Carlin, Salvacion Gatchalian, Deborah Lehmann, Lulu Muhe, E Kim Mulholland, and WHO Young Infants Study Group. Predictors of neonatal sepsis in developing countries. 22:711–717, 08 2003.

[23] Niccolo Pignatelli. Design of a mobile kit for cardiovascular disease screening in resource constrained environments. Master's thesis, Massachusetts Institute of Technology, 2017.

[24] National Rural Health Mission of India. *Induction Training Module for ASHAs*. National Rural Health Mission of India. (Accessed on 05/23/2018).

[25] Who | the who child growth standards. (Accessed on 05/24/2018).

[26] Merriam-Webster. "physiology.". (Accessed on 05/24/2018).

[27] Suchi Saria, Anand K. Rajani, Jeffrey Gould, Daphne Koller, and Anna A. Penn. Integration of early physiological responses predicts later illness severity in preterm infants. *Science Translational Medicine*, 2(48):48ra65–48ra65, 2010.

[28] Hao-Yu Wu, Michael Rubinstein, Eugene Shih, John Guttag, Frédo Durand, and William T. Freeman. Eulerian video magnification for revealing subtle changes in the world. *ACM Trans. Graph. (Proceedings SIGGRAPH 2012)*, 31(4), 2012.

[29] Ennio Gambi, Angela Agostinelli, Alberto Belli, Laura Burattini, Enea Cippitelli, Sandro Fioretti, Paola Pierleoni, Manola Ricciuti, Agnese Sbrollini, and Susanna Spinsante. Heart rate detection using microsoft kinect: Validation and comparison to wearable devices. *Sensors*, 17(8), 2017.

[30] E. J. Wang, W. Li, J. Zhu, R. Rana, and S. N. Patel. Noninvasive hemoglobin measurement using unmodified smartphone camera and white flash. In *2017 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 2333–2336, July 2017.

[31] Edward Jay Wang, William Li, Doug Hawkins, Terry Gernsheimer, Colette Norby-Slycord, and Shwetak N. Patel. Hemaapp: Noninvasive blood screening of hemoglobin using smartphone cameras. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '16, pages 593–604, New York, NY, USA, 2016. ACM.

[32] Edward Jay Wang, Junyi Zhu, Mohit Jain, Tien-Jui Lee, Elliot Saba, Lama Nachman, and Shwetak N. Patel. Seismo: Blood pressure monitoring using built-in smartphone accelerometer and camera. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pages 425:1–425:9, New York, NY, USA, 2018. ACM.

[33] Alex Mariakakis, Megan A Banks, Lauren Phillipi, Lei Yu, James Taylor, and Shwetak N Patel. Biliscreen: Smartphone-based scleral jaundice monitoring for liver and pancreatic disorders. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(2):20, 2017.

[34] Lilian de Greef, Mayank Goel, Min Joon Seo, Eric C. Larson, James W. Stout, James A. Taylor, and Shwetak N. Patel. Bilicam: Using mobile phones to monitor newborn jaundice. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '14, pages 331–342, New York, NY, USA, 2014. ACM.

[35] Alex Mariakakis, Jacob Baudin, Eric Whitmire, Vardhman Mehta, Megan A Banks, Anthony Law, Lynn McGrath, and Shwetak N Patel. Pupilscreen: Using smartphones to assess traumatic brain injury. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 1(3):81, 2017.

[36] A Mariakakis, E Wang, S Patel, and JC Wen. A smartphone-based system for assessing intraocular pressure. In *Engineering in Medicine and Biology Society (EMBC), 2016 IEEE 38th Annual International Conference of the*, pages 4353–4356. IEEE, 2016.

[37] K. Lunze and D. H. Hamer. Thermal protection of the newborn in resource-limited environments. *Journal Of Perinatology*, 32:317 EP –, Mar 2012. State-of-the-Art.

[38] Tavpritesh Sethi, Aditya Nagori, Ambika Bhatnagar, Priyanka Gupta, Richard Fletcher, and Rakesh Lodha. Validating the tele-diagnostic potential of affordable thermography in a big-data data-enabled icu. In *Proceedings of the Special Collection on eGovernment Innovations in India*, ICEGOV '17, pages 64–69, New York, NY, USA, 2017. ACM.

[39] Biao Zhang. Decision analysis of novel point-of-care diagnostics for Pediatric Pneumonia : implementation in Developing countries with tiered healthcare systems. Master's thesis, Massachusetts Institute of Technology, 2015.

[40] Linda T. Wang, Robert H. Cleveland, William Binder, Robert G. Zwerdling, Caterina Stamoulis, Thomas Ptak, Mindy Sherman, Kenan Haver, Pallavi Sagar, and Patricia Hibberd. Similarity of chest x-ray and thermal imaging of focal pneumonia: a randomised proof of concept study at a large urban teaching hospital. *BMJ Open*, 2018.

[41] Carina Barbosa Pereira, Xinchi Yu, Michael Czaplik, Rolf Rossaint, Vladimir Blazek, and Steffen Leonhardt. Remote monitoring of breathing dynamics using infrared thermography. *Biomed. Opt. Express*, 6(11):4378–4394, Nov 2015.

[42] Opencv: Cascade classifier training. (Accessed on 05/24/2018).

[43] Paul Viola and Michael Jones. Rapid object detection using a boosted cascade of simple features. pages 511–518, 2001.

[44] Paul Viola and Michael J. Jones. Robust real-time face detection. *Int. J. Comput. Vision*, 57(2):137–154, May 2004.

[45] Daniel Chamberlain. Design and validation of mobile kit and machine learning algorithms for pulmonary disease screening and diagnosis. Master's thesis, Massachusetts Institute of Technology, 2017.

[46] Fadel Adib, Zachary Kabelac, Dina Katabi, and Robert C. Miller. 3d tracking via body radio reflections. In *Proceedings of the 11th USENIX Conference on Networked Systems Design and Implementation*, NSDI'14, pages 317–329, Berkeley, CA, USA, 2014. USENIX Association.