UNIVERSITY OF CAPE TOWN

FACULTY OF THE HEALTH SCIENCES

DIVISION OF BIOMEDICAL ENGINEERING

Mobile phone – based evaluation of latent tuberculosis infection

MASTERS FULL DISSERTATION

submitted to

THE UNIVERSITY OF CAPE TOWN

in fulfilment of the requirements for the degree of

MASTER OF SCIENCE IN BIOMEDICAL ENGINEERING

(MSc BME)

Author: *Safa Naraghi – NRGSAF001*

Supervisor: *Dr Tinashe Mutsvangwa*

Co-Supervisor: *Professor Tania Douglas*

Date: 31 October 2017

DECLARATION OF AUTHORSHIP

I, Safa Kagiso Naraghi, hereby declare that the work on which this dissertation/thesis is based is my original work (except where acknowledgements indicate otherwise) and that neither the whole work nor any part of it has been, is being, or is to be submitted for another degree in this or any other university.

I empower the university to reproduce for the purpose of research either the whole or any portion of the contents in any manner whatsoever.

Signature: signature removed

Date: 31/10/2017

ABSTRACT

UNIVERSITY OF CAPE TOWN

Masters Dissertation

Mobile phone-based evaluation of latent tuberculosis infection

by Safa Naraghi

NRGSAF001

The tuberculin skin test (TST) is the most widely used method for detecting latent tuberculosis (TB) infection (LTBI) in adults and active TB disease in children. This work presents the development of a screening tool to detect LTBI's, which works in conjunction with the TST and serves as an alternative for measuring the TST induration. The screening tool makes use of a mobile application developed on the Android platform to capture images of an induration, and photogrammetric reconstruction using Agisoft PhotoScan to reconstruct the induration in 3D, followed by 3D measurement of the induration with the aid of Python functions. The screening accuracy of the developed process was tested using a 3D printed induration and an HTC One smartphone to capture images. In this accuracy test, the developed screening tool was found to measure indurations more accurately than current measurement methods, as indicated by the lower standard deviation produced. An experiment to simulate real-world conditions was conducted by using the developed screening tool on a set of mock skin indurations, created by a make-up artist, and evaluating its performance. It was found that the height of the skin induration and definition of its margins are the most significant factors that influence the accuracy of the screening tool under simulated real-world conditions. Future work should explore possible improvements to the developed image capture protocol and the bimodal segmentation methods employed in this project.

ACKNOWLEDGMENTS

I would like to express my sincere thanks and gratitude to the following persons:

My supervisor and co-supervisor, Dr Tinashe Mutsvangwa and Professor Tania Douglas, respectively, for providing guidance over the course of producing this dissertation. Without their continuous and selfless support, I would not have been able to complete this dissertation.

Dr Molebogeng Rangaka, epidemiologist and researcher at University College London, for proposing this dissertation topic and providing continual guidance with regard to the context of the problem we aimed to tackle.

Dr Rene Goliath, clinician and researcher at the University of Cape Town, for providing clinical advice and assisting with the execution of the simulated clinical experiments.

The 2015, 2016 and 2017 Biomedical Engineering MSc student cohorts, for all the advice and help over the last two years. Their insights were invaluable and certainly guided me through the toughest moments in completing this project.

My colleagues in the Medical Imaging Research Unit and the Lodox office, for their honest critique of my work, their accommodating attitude during the period of my experiments and the laughter and joy they brought to my work environment.

The National Research Foundation (NRF), which supported the work through an Innovation Scholarship, and the DST/NRF South African Research Chairs Initiative, which provided project support through the Research Chair in Biomedical Engineering & Innovation.

My brothers, Vafa and Sama Naraghi, for the countless distractions during the making of this dissertation, without which I would have completed my dissertation significantly earlier.

My parents, Sefidvash and Marina Naraghi, for their advice, love and unwavering support through my studies at UCT.

Ya Baha'u'l-Abhá

.

TABLE OF CONTENTS

.

LIST OF ILLUSTRATIONS

.

ABBREVIATIONS

AIDS **A**cquired **I**mmunodeficiency **S**yndrome

APS **A**gisoft **P**hoto**S**can

HIV **H**uman **I**mmunodeficiency **V**irus

IGRA Interferon-Gamma Release Assays

IMU **I**nertial **M**easurement **U**nit

LED **L**ight **E**mitting **D**iode

LTBI **L**atent **T**uberculosis **I**nfection

MSE **M**ean **S**quared **E**rror

SIFT **S**cale-**I**nvariant **F**eature **T**ransform

SFM **S**tructure-**F**rom-**M**otion

SSIM **S**tructural **S**imilarity **I**ndex

TB **T**uberculosis

TST **T**uberculin **S**kin **T**est

WHO **W**orld **H**ealth **O**rganisation

1. INTRODUCTION

1.1 Background

The tuberculin skin test (TST), is the most widely used method for detecting latent tuberculosis (TB) infection (LTBI) in adults and active TB disease in children (WHO, 2015). LTBI is characterised by infection with *Mycobacterium tuberculosis* without symptoms of active disease. The TST is carried out by injecting tuberculin into the patient's arm which may result in the formation of an induration (in both infected and uninfected persons) on the skin after 48 to 72 hours. The diameter of the induration measured across the forearm is used to infer LTBI. The risk of subsequent TB in persons with positive reactions is 2-fold higher than those with negative reactions (Rangaka et al., 2012). The World Health Organization (WHO) cites that 30 of the 68 countries that implement a systematic method of testing and treatment of LTBI solely rely on the TST for diagnosis; with the other 38 countries using a combination of the TST and other diagnostic tests (WHO, 2015). With 28% of the world's tuberculosis cases in 2014, Africa is the most heavily burdened by the disease relative to the population with more than double the global average of tuberculosis cases per 100 000 people (WHO, 2015).

Major disadvantages of the TST include the requirement for clinical follow-up and the highly operator-dependent nature of the test (Haghdoost et al., 2014). Failure to attend follow-up within the evaluation period invalidates test results. People with untreated LTBI are at risk of developing active TB and transmitting infection in the community. Studies have shown that widely used methods of measuring TST indurations are inaccurate with high intra-observer standard deviations (Meyer et al., 1951). Moreover, persons infected with latent TB could miss the opportunity to receive treatment. Further obstacles to returning include time, money, and distance to the healthcare facilities (Oliveira et al., 2011).

Africa is considered the least developed continent in the world with just over 60% of its population living outside urban areas having poor access to health facilities (UN, 2014). Bridging this gap of disconnectedness, however, is the rising use of the smartphone in Africa. Africa, in line with other developing countries, has seen exponential growth in the use of smartphones amongst its population (ITU, 2015), with internet access on the

continent following similar trends (Gillwald et al., 2013). The combination of the rise of the use of mobile devices in Africa, increased access to the internet and the burden of tuberculosis on the continent provides an opportunity for the implementation of effective mobile health screening tools. Mobile health screening is the use of mobile devices to screen medical conditions, an innovation that could bridge the existing gap of reliable constant access to health care professionals in remote areas of Africa (Wilson et al., 2009). The use of smartphones to transfer medical information to a central location for expert analysis may provide a new channel for Africa's rural population to access health care. Mobile-health solutions have been explored for diagnosing active TB but not for LTBI detection (Denkinger et al., 2013). Pertaining to the TST particularly, mobile health screening has the potential to alleviate the need to visit a health professional for diagnosis 48 to 72 hours after the test is performed.

1.2 Aim and objectives

The aim of this project was to develop an automated system, procedure and protocol for screening latent TB using images of a TST induration. This aim was achieved through the following objectives:

1) Establishing a protocol for acquiring suitable images of an induration using a smartphone.
2) Developing a smart phone application to implement the developed image capture protocol.
3) Creating a 3D reconstruction of the induration through an automated process using the structure-from-motion (SFM) method implemented by an off-the shelf 3D modelling program.
4) Developing an algorithm which extracts information from the 3D reconstruction, applies relevant processing methods, and enables automated measurement of the critical diagnostic diameter of the induration.
5) Evaluating the developed solution and comparing it with a current measurement method.

1.3 Scope and limitations

This project was primarily focused on developing a tuberculosis screening tool for use on a low-cost smartphone. Although the entire screening system could be implemented anywhere in the world, the design of the system will be focused on use in resource-limited

settings as these locations are where the problems associated with follow-up TST adherence are most prominent and where the use of the lowest quality/cost smartphones is more likely. Emphasis was placed on developing a methodology and procedure for acquiring images in a user-friendly and efficient manner while also ensuring an accurate final screening result. The 3D reconstruction of the induration was implemented using Agisoft PhotoScan (Agisoft LLC, Russia: www.agisoft.com), a commercially available tool for photogrammetric 3D reconstruction and real-world measurement. Python (Python Software Foundation: www.python.org) scripting was used as an interface to Agisoft PhotoScan and for all image processing used in this project. The application was tested using an HTC One M8 (www.htc.com) smartphone (5MP primary camera, 1.4µm pixel size, 1/2.5-inch sensor size and a f/1.7, 26mm lens) on the Android platform.

1.4 Structure of dissertation

Chapter 2 presents a review of the literature pertaining to tuberculosis, the mobile health landscape and relevant 3D reconstruction methods. Chapter 3 outlines the objectives, scope and limitations of the project along with a general overview of the developed solution with an emphasis on how various parts of the system were connected. Chapter 4 presents the methodology used in developing the image acquisition protocol and Chapter 5 outlines the functions of the developed application which assist in implementing the aforementioned image acquisition protocol. Chapter 6 presents an overview of the photogrammetric 3D reconstruction process used for this project and the implemented database management system. Details of the algorithm used to evaluate the results from the 3D reconstruction are presented in Chapter 7. Testing of the application and entire system on both test objects and simulated skin indurations is presented and discussed in Chapter 8. Finally, a general discussion with conclusions and recommendations for future work appear in Chapter 9.

2. LITERATURE REVIEW

2.1 Tuberculosis and Tuberculosis testing methods

2.1.1 Tuberculosis pandemic

In 2012, it was reported that TB killed 1.3 million people across the world (WHO, 2013); this figure has since risen to 1.5 million in 2014 (WHO, 2015). This places TB along with HIV as the leading causes of death from a single infectious disease worldwide (Ortblad et al., 2013). Not only does the disease threaten the lives of millions across the world, it also affects the economic stability of households. Lopez et al. (2006) reported that TB is the leading cause of death by disease amongst economically active age groups globally.

Due to the HIV epidemic, annual TB notification rates in sub-Saharan African countries increased almost three-fold in 2005 (WHO, 2006). The two diseases tend to coincide in one geographical area as the pathogenesis of TB is changed by HIV and thus the risk of TB infection is increased (Cahn et al., 2003). Latent tuberculosis infection is characterised by patients being infected with Mycobacterium Tuberculosis (M.tuberculosis) without having any common symptoms of active tuberculosis. With 2-10% of HIV infected individuals developing active TB from latent TB annually, HIV is seen as a principal risk factor of latent M.tuberculosis activation and spread of the disease (Chaisson et al., 1996). The presence of active TB in HIV/AIDS infected individuals significantly decreases the chances of survival and thus early detection of LTBI's and prevention of LTBI activation to active TB is imperative (Abdool Karim et al., 2011). In addition, with reference to non-HIV infected individuals, the risk of subsequent TB in persons with positive LTBI diagnosis is 2-fold higher than those with negative diagnosis (Rangaka et al., 2012). While TB preventative therapy is widely accessible, it is reliant on accurate and widespread detection of latent TB (WHO, 2015).

In a review of methods of LTBI diagnosis, Mardani and Abtahian (2015) claim that accurate latent TB diagnosis is one of the greatest challenges facing the fight against the TB pandemic. They also noted that until the beginning of the 20th century, the tuberculin skin test (TST) (see description below) was the only method of detecting latent TB in patients. Since then, new methods and procedures for detecting latent TB have been developed and

implemented (Mardani and Abtahian, 2015). Two current methods used to diagnose latent TB are detailed below.

2.1.1.1 Tuberculin skin test

The TST relies on the hypersensitivity response of M.tuberculosis and tuberculin and can be used to diagnose LTBI's in adults and active TB in children. The tuberculin is injected into the skin of the patient (usually in the forearm region) and, in possibly infected patients, results in the formation of an induration due to the reaction between M.tuberculosis and tuberculin. Based on the size of the resulting induration, the presence of latent TB is determined. Mardani and Abtahian (2015) outlined the following limitations of the TST:

- TST sensitivity is reduced by the presence of certain immunodeficiency diseases such as HIV.
- A decrease in TST specificity might arise in patients that have received Bacille Camette-Gurin (BCG) vaccination after infancy (Kritzinger et al., 2009).
- Two health care visits are required to complete the test. The first, for the injection of the tuberculin and the second, for evaluation of the induration. Linas et al. (2011) estimated that the effect of the requirement to have two health care visits is a 10% loss of results from the test.

With reference to the effect of HIV on the appearance of the induration, a clinician at the University of Cape Town (UCT) was consulted on possible variations. It was noted that the induration usually presents in a form resembling an ellipse and is less prominent (shallow indurations and unclear induration margin definition) for HIV infected individuals (Goliath, personal communication, 2017). The TST has its advantages in the large amount of clinical research and experience supporting it as it has been studied and used extensively during the past century (Mardani and Abtahian, 2015). For this reason, along with its low cost, the TST is still widely used in developing countries in Sub-Saharan Africa (WHO, 2015).

2.1.1.2 Interferon Gamma Release Assay (IGRA)

IGRA makes use of measuring the immune response to TB proteins in the blood to determine if an individual is infected with tuberculosis (Mardani and Abtahian, 2015). By mixing antigens derived from M.tuberculosis with peptides and exposing it to the blood, a release of interferon gamma (INF-γ) will be present in an individual infected with

tuberculosis (Trajman et al., 2013). Mardani and Abtahian (2015) outlined the following limitations of the IGRA:

- The blood sample should be processed within 8-30 hours of being collected. This has a significant impact on populations in remote areas which are distant from a processing centre.
- Children younger than the age of five cannot be tested using this method.
- The test is not accurate for people recently exposed to TB or immunocompromised individuals.

The advantages of IGRA are the following: the test can be conducted in a single visit with a clinician; the results of the test can be available within a 24 hour period; and the test is unaffected by BCG vaccination and most environmental mycobacteria (Mardani and Abtahian, 2015).

2.1.1.3 Administration of the tuberculin skin test

The following procedure for conducting the TST was published on a wall-chart for clinicians by the Centers for Disease Control and Prevention (CDC, 2013). The procedure is separated in to 3 main sections (administration (A), reading (B) and interpretation (C)) with subsections pertaining to each section.

A. Administration

Prior to administering the tuberculin, the clinician conducts a risk assessment which takes into consideration factors such as recent exposure to TB and clinical conditions which could possibly increase the risk for TB infection. As shown in Figure 2.1, the clinician locates and disinfects an area about 10 cm distal to the elbow joint which is free of any obstructions (i.e. scars).

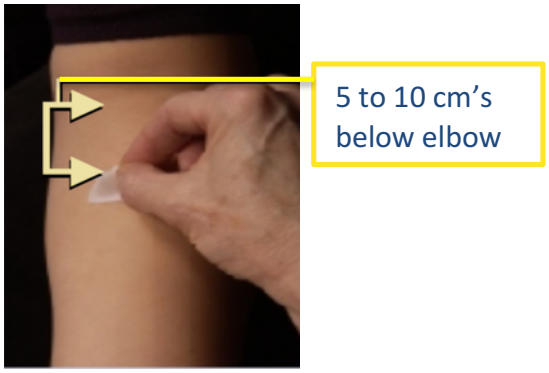

Figure 2.1: Location of injection for TST (image from CDC (2013)).

Using a syringe, the clinician injects about 0.1 mm of tuberculin under the skin of the patient's forearm in the selected area. This results in the formation of a wheal 6 to 10 mm in diameter (Figure 2.2).

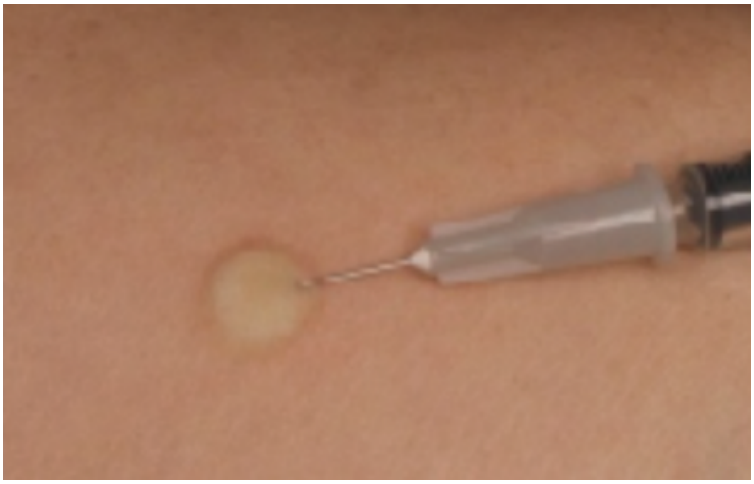

Figure 2.2: Pale wheal on skin after TST injection (image from CDC (2013)).

B. Reading

Once the administration is complete, the patient may leave and is asked to return within 48 to 72 hours. Upon returning, the clinician identifies the induration formed (if any) on the patient's skin. As shown in Figure 2.3, the induration may be surrounded by erythema (reddening due to irritation of the skin).

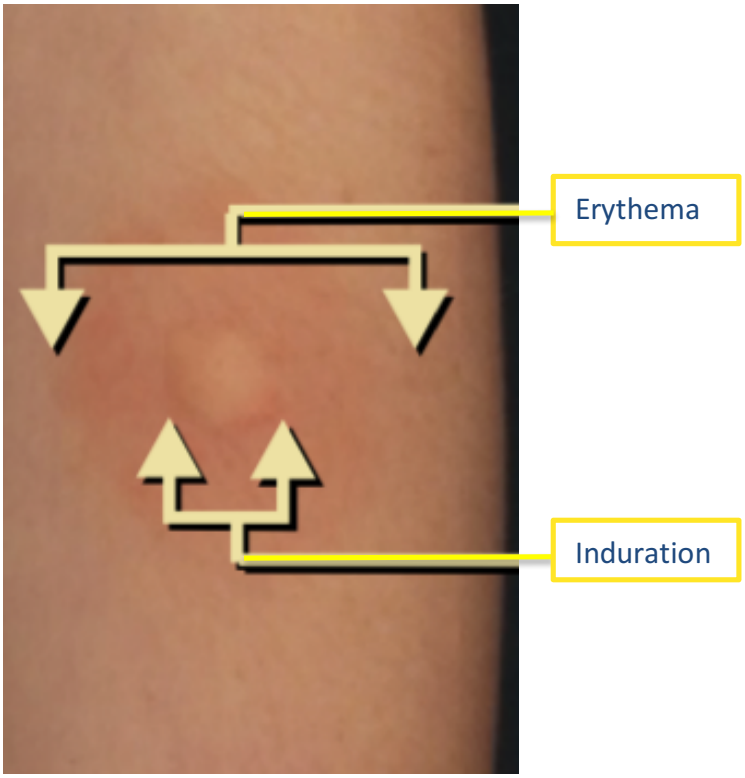

Figure 2.3: TST induration and erythema on the skin after reaction (image from CDC (2013))

The erythema may distort the edges of the induration. As such, it is important for the clinician to identify the margins of the induration by palpating using his/her fingertips or by using ballpoint pen method (outlined below) to identify the edges of the induration. As shown in Figure 2.4, the widest edges of the induration across the forearm are marked with a pen.

Figure 2.4: Marking of the TST induration edges (image from CDC (2013)).

Using a ruler, the clinician measures the diameter of the induration across the forearm perpendicular to the long axis of the arm (Figure 2.5). If the measurement falls between graduation lines of the ruler, the lower graduation is used as the measurement.

Figure 2.5: Measurement of the TST induration using ruler (image from CDC (2013))

C. Interpretation

The size (diameter) of the induration, categorised into three pre-determined ranges, gives an indication of whether the TB mycobacteria is present based on patient-specific environmental and health conditions. These size ranges are 0 mm to 5 mm, 5 mm to 10 mm and 10 mm to 15 mm. Table 2.1 shows a summary of patients' health conditions and environments which correspond to distinctive measurement bins for a positive diagnosis of LTBI's.

Table 2.1: Table showing induration diameter for TB diagnosis in various patient-specific circumstances (adapted from CDC (2013)).

Environment/Condition	Induration Diameter		
	≥ 5 mm	≥ 10 mm	≥ 15 mm
Persons infected with HIV	✓	✓	✓
Recent contact with TB case patient	✓	✓	✓
Persons with indications of prior TB infection	✓	✓	✓
Patients with organ transplants	✓	✓	✓
Recent immigrants from countries with high TB prevalence	✗	✓	✓
Injection drug users	✗	✓	✓
Employees/residence of high-risk congregate settings (i.e. prisons, homeless shelters)	✗	✓	✓
Mycobacterial laboratory personnel	✗	✓	✓
Persons with silicosis or diabetes mellitus	✗	✓	✓
Children under 5 years of age	✗	✓	✓
Infants, children and adolescents exposed to adults at high risk of developing active TB	✗	✓	✓
Persons with no known risk factors for TB	✗	✗	✓
LEGEND: ✓ Positive TB diagnosis ✗ Negative TB diagnosis			

2.1.1.4 Possible errors in reading of TST result

When administering the TST, the measure of skin reaction due to tuberculin administration is assessed often without taking into consideration variations and unavoidable errors in administering tests and reading test results (Meyer et al., 1951). An ideal administering and measuring technique would result in accurate diagnosis, however, this cannot be truly achieved in practice. Various small errors such as the depth of the administered injection, induration boundary definition and quantity of tuberculin injected may skew the perceived

result of the test (Meyer et al., 1951). Huebner et al. (1993) argue that only a small percentage of errors in diagnosis are as a result of administration and reading errors. However, Pouchot et al. (1997) and Meyer et al. (1951) have shown that considerable errors and misdiagnosis occur due to these factors. There are various contributors to errors in diagnosing LTBI's using the TST, however, this section will focus on the error that occurs as a product of reading the result of the TST only.

Variation in the measurement of the same induration is common and in a study comprising 6 large test groups with a total of 3600 participants, Meyer et al. (1951) found that the reading error of the test result (diameter of induration) had a standard deviation of between 1.09 mm and 1.75 mm. The authors concluded that, in the worst case, 95% of indurations which have a true diameter of 10 mm would fall in the range of 6.5 to 13.5 mm and approximately 5 of 100 would fall either below 6.5 mm or above 13.5 mm in diameter. Considering that the margins for diagnosing LTBI's according to this test require millimetre accuracy, these results could be significant sources of misdiagnosis. Furthermore, experienced clinicians were used in the abovementioned study and the observed error may have been larger if less experienced clinicians were used.

Pouchot et al. (1997) reported that reading the resulting diameter of the induration was one of the most important sources of error associated with the TST. In this study, two experienced clinicians measured the same induration multiple times. The results showed that 95% of the intra-observer errors fell between -2.68 mm and +2.96 resulting in a 12% diagnosis reclassification rate in the test group. The authors also noted that error in reading the result can be further compounded by an error in marking the edges of the induration. They attributed the error in marking to be predominantly a result of two factors: location of the induration edge furthest from the long axis of the arm and location of the start (raised section of skin) of the induration.

Another source of possible human error cited by Huebner et al. (1993) is conscious or unconscious bias on the part of the reader. This bias could be as a result of an array of factors such as awareness of cut-off values, knowledge of the amount of tuberculin administered and possible faults during administration (Pouchot et al., 1997). Furthermore, Pouchot et al. (1997) noted that inexperienced test interpreters or undertrained clinicians

may contribute to errors in correctly assessing the result of the TST. This could be of significance to remote resource-limited areas where there is a shortage of well-trained medical staff.

2.1.1.5 TST patient experience

Although there is limited academic research pertaining to the patient experience of TST evaluation, there are indications that the test may be uncomfortable (Goliath, personal communication, 2017). A common method of marking the edges of the induration is termed the "ballpoint pen method", where the point of a pen is placed on the forearm and moved toward the centre point of the induration (Huebner et al., 1993). When resistance is felt, due to reaching the edge of the induration, the pen is lifted and the marking of the induration boundary is formed. Since erythema (reddening and irritation of skin) is commonly found around the induration this procedure may be uncomfortable for the patient. Furthermore, as noted in the procedure of the TST, there is significant use of the fingertips in applying pressure to the skin of the forearm, which could be a source of further discomfort for the patient.

2.1.1.6 Comparison of induration measurement methods

Although it is not common to use instruments other than a ruler to measure the diameter of the TST induration, Morán-Mendoza et al. (2013) compared the use of a ruler and Vernier calliper. The study compared the use of 4 methods:

1) Palpation and ruler method.
2) Ballpoint pen and ruler method.
3) Palpation and Vernier calliper method.
4) Ballpoint pen and Vernier caliper method.

During the experiment, the TST was administered on both arms in 78 subjects. The "ruler and ballpoint pen" and "ruler and palpation" methods had the least variability in recorded measurements. Alongside assessing measurement methods for evaluating the TST, the study also investigated the best time (post administration) to measure the result of the test. It was concluded that the result should ideally be assessed 72 hours after administration with either the "ruler and ballpoint pen" or "ruler and palpation" methods.

2.1.1.7 TST follow-up visit limitations

A common issue with completing the TST is the adherence to follow-up visits, 48 to 72 hours after administration, for assessment of the result of the test. This results in a loss of diagnostic data as the test cannot be read for diagnostic purposes beyond 72 hours post-administration. Non-returners may have latent TB and would not be aware of it due to incompletion of the TST follow-up assessment. The estimated 10% loss in test results presented above was deduced from data pertaining to the USA. However, Haghdoost et al. (2014) have shown that this problem is magnified in developing regions. In their study on 6403 children in Iran, close to 1000 children did not return for evaluation of the TST in the recommended period. The authors noted that most of the children who did not return were from remote areas and transport, money and time constraints may have played a role in the decline in follow-up visits.

Studies conducted in rural clinics in Africa show the same trend. Aisu et al. (1995) performed a study on operational assessment of HIV-associated TB in clinics in Uganda. Initially, 1524 patients were selected to participate in the study, however, this number decreased by 200 for various reasons expressed by the patients. One of these reasons was that the clinic was too far from where some patients lived. Of the 1344 participants that took the TST, 1094 (81.4%) returned for reading of the test results. Amongst factors associated with lower income populations (such as monetary constrains and access to transportation), the authors noted that participants might have been discouraged to return as the clinician only worked part-time at the clinic, which imposed constraints on the times of the day for the follow-up visit.

2.1.2 Summary of Tuberculosis and TB testing methods

It is imperative that low and medium income countries that have high TB-incidence and HIV rates diagnose latent TB and implement active preventative treatment. A possible obstacle to investment in diagnosis tools for latent TB in such countries is the ease of use of the diagnostic test and possible barriers to completion of the test such as equipment or specialised personnel (Mardani and Abtahian, 2015). The low cost of TST makes it the preferred method for diagnosing TB in developing countries. Two main groups rely heavily on the use of TST for latent TB detection: adults in developing regions and children under

the age of 5 (for testing of LTBI and active TB). For adults in developing regions, the TST is the most accessible test given the low cost of the test. Children under the age of five are not eligible to use the IGRA test and thus rely on the TST for latent and active TB detection.

One of the major barriers to the completion of the TST in developing countries is the lack of skilled personnel available to assess the result of the test and contributes to the other major barriers to completing the TST such as time, money and distance. This is because the few trained personnel available are spread over large geographical areas and are often costly and time-consuming to access for the entire population of a region.

Finally, the current methods used for evaluation of the TST have a large intra-observer variation (measurement standard deviation greater than 1.1 mm) which results in misdiagnosis in some cases. Clinicians may benefit from methods with better precision.

2.2 Mobile healthcare and internet access in Africa

2.2.1 Mobile technology and internet access in Africa

With an estimated mobile phone subscription annual growth rate of 65%, Africa is the leading continent in mobile phone growth rates (Hosman and Fife, 2012). The World Bank estimated that there are 650 million mobile phone users in Africa. Although the bulk of mobile phone users in Africa reside in urban areas of the continent, there has been an expansion in mobile phone use in remote and rural areas. Swahn et al. (2014) reported that, in a survey of youth in rural areas in sub-Saharan Africa, 46.9% of the participants reported owning a mobile phone which they used to access the internet. This provides an avenue for accessing many remote populations in the region. Alongside the rise of access to mobile internet in Africa is the surge in access to general internet use. Amongst populations that did not own internet enabled mobile phones, Gillwald et al. (2013) found that there was an increase of up to 33.7% of people who had gained access to the internet in one form or another over a 4 year period (2009-2012).

2.2.2 Mobile healthcare in Africa

Mobile health, commonly known as **mHealth**, is an emerging field in healthcare (Swahn et al., 2014). It was first defined broadly as “emerging mobile communications and network technologies for healthcare”(Istepanian et al., 2006). MHealth makes use of various mobile

or monitoring devices to provide health care information or screening/diagnostic services (Folaranmi, 2014).

Innovation around mHealth is on a sharp rise in Africa with countries like Ethiopia, Nigeria, Kenya and South Africa leading the way. Folaranmi (2014) reported that the innovation in this industry is mainly spurred on by factors such as the rising mobile network penetration in remote areas of the continent, low cost of mobile devices, and accessible technologies that make use of mobile applications for the delivery of health services. Furthermore, the implementation of low cost and flexible data plans on the continent is expected to expand this industry further (Folaranmi, 2014).

The benefits of mHealth are seen in both the monetary and community upliftment spheres. Collecting data using a mobile phone could provide a cost saving avenue for government by reducing the amount of paper used and data entry costs (Folaranmi, 2014). Furthermore, Folaranmi (2014) reported that the implementation of mHealth programs could also reduce the cost of travel for patients and doctors. In the community upliftment sphere, they reported that the use of a web based mHealth data collection tool in Uganda recently proved to be successful in kerbing disease outbreaks and has been implemented in 5000 health facilities around the country. Below are some examples of mobile healthcare applications based on image screening.

Mobile phone based clinical microscopy

Light microscopy can be used in various diagnostic applications of infectious diseases. In an effort to overcome limited accessibility and widen the use of light microscopy, Breslauer et al. (2009) developed a portable mobile phone-mounted light microscope. Amongst other tests, they demonstrated its potential by imaging sputum infected with M.tuberculosis in flourescence light with LED excitation. Using the developed device, images of sputum samples were captured and sent to an expert for analysis. The authors concluded that the 3.2 megapixel images captured using the device were sufficient for easy identification of TB bacteria in the sample. They also noted that the increasing resolution of smart phone cameras and computing ability could increase the quality of their results and open other avenues for TB screening.

Ear biometrics smartphone application

In Bargal et al. (2015), it was suggested that ear biometrics could provide a solution to the problem of medical record management in clinics in less developed countries where patients may not possess identification documents. In the absence of reliable patient identification, consequences such as repeated treatments and disruption of continuous patient care are prevalent (Bargal et al., 2015). The authors addressed this problem by means of a mobile application to be used by clinicians to image patient's ears and send this image to a central processing centre for identification.

The application was developed as a tool to capture images which were then processed and matched to other images of ears in a database. To apply the matching process, the Scale Invariant Feature Transform (SIFT) method was implemented in Matlab at a central processing site. Prior to using the SIFT method, the image was converted to grey scale, cropped and down sampled. It was found that implementing such image processing techniques reduced the variability of the environment in which image was taken in, thus making the process of feature recognition more accurate. Furthermore, image acquisition protocols such as a box alignment requirement (aligning the ear in a box which is visible on the image capture screen) to facilitate easier feature identification, were implemented. It was also found that using the flash function on mobile devices when taking the image increased the accuracy of the process as this mimicked standard lighting conditions.

SkinVision™

SkinVision (Skin Vision BV., Netherlands) is a smartphone application which was developed as a means of diagnosing and self-monitoring pigmented moles on the skin which could possibly be diagnosed as malignant melanoma (Maier et al., 2015). This application assists patients in flagging risks of melanoma by taking a picture of the mole on the skin which is then sent for processing. In contrast to other similar smartphone applications on the market which are based on image processing applications established through consultation of a dermatologist, SkinVision uses a risk assessment algorithm based on Fractal theory (Losa, 2011). Fractal methods are used to quantify and describe geometric features, shapes, and patterns. Such methods have also been used for texture classification in microscopic medical applications (Losa, 2011). The algorithm uses a variety of factors such as the

number of regions with different textures and the size of the regions (as identified by region boarders) to classify the risk associated with a lesion (Maier et al., 2015).

In tests conducted by Maier et al. (2015), it was found that although evaluation conducted by a trained dermatologist yielded better results, the smartphone application detected melanoma with an accuracy of 81% and a sensitivity and specificity of 73% and 83% respectively. There is room for the algorithm to be improved to make it a robust diagnostic tool, however, the application is currently best suited as a screening tool to give a basis for patient-doctor interaction.

Smartphone application for assessment of burn injury

Wallis et al. (2016) developed a smartphone application and cloud server for analysis of skin wounds of burn victim patients. The application was developed for use in resource-poor settings where outcomes of patients with severe burn injuries have worsened in recent years. The application comprises 4 major components/protocols: image acquisition, total body surface calculator, cloud based image storage and expert analysis. After gathering information from the patient using the application interface by means of a short survey, the pictures of the wound are captured by the patient or by an on-hand clinician, following a set image acquisition protocol. The user is then prompted to colour the region of the wound an avatar generated by the application and estimate the wound depth. All information pertaining to the wound (including pictures) are sent to an expert for analysis and recommendation for treatment. Although the application does not use automated techniques for diagnosis, it provides a link between experts and patients in remote areas that did not exist before. Clinical trials of the application have not been conducted yet and therefore the impact of the application has not been tested.

2.2.3 Summary of mobile healthcare and internet access

Access to the internet in some form or another is becoming widely available to populations in resource-limited regions. However, there is little indication in the literature of the speeds and cost of accessing the internet in these regions. Being cognisant of data speed and cost in developing mHealth applications for resource-limited settings will ensure greater access and adoption.

Drawing from the abovementioned mHealth solutions, the industry standard is to develop screening tools rather than diagnostic tools until such a time that the technologies mature. The use of a centrally located machine for processing and data storage as in the burn injury and ear biometrics applications are important considerations when using smartphones with low processing power and allow for increasing the accessibility of the developed mobile solution.

2.3 Mobile phone-based photogrammetry

2.3.1 Photogrammetry

In its simplest form, photogrammetry is defined as the science of using photographs to take real-world measurements (Walford, 2007). Schenk (2005) defined it more technically as "the science of obtaining reliable information about the properties of surfaces and objects without physical contact with the objects, and of measuring and interpreting this information". There are various outputs of photogrammetry such as maps or 3D models of objects that exist in the real world (Walford, 2007).

Walford (2007) explained and expanded on two types of photogrammetry: aerial photogrammetry and close-range photogrammetry. Aerial photogrammetry usually involves a camera mounted on an aircraft which is directed vertically downward to face the ground. The processed output is commonly a map of the terrain below the aircraft, however, more modern methods are currently used for map generation. In contrast, they described close-range photogrammetry as a process which makes use of images taken close to the subject of interest. These images are taken using a handheld device or a camera mounted on a tripod stand. The output of this process is usually 3D models and measurements of various object geometries in the 3D realm. Walford (2007) outlined a variety of applications of close-range photogrammetry which range across the modelling and measurement of forensic information, archaeological artefacts, and engineering structures and mines.

2.3.2 Structure-from-motion close range photogrammetry

Close-range photogrammetry is mainly concerned with the geometric reconstruction of features and may be implemented using structure-from-motion (SFM). As the name suggests, SFM aims to estimate the 3D structure of a static scene from a series of still frames (images) captured in some motion (in 3D space) in the scene (Bolles et al., 1987).

Utilising information about the motion of the camera, the SFM algorithm analyses slices of the object of interest and forms a 3D reconstruction of the object. Figure 2.6 illustrates how a set of 3 images of an object (P) are captured in a linear motion for SFM reconstruction.

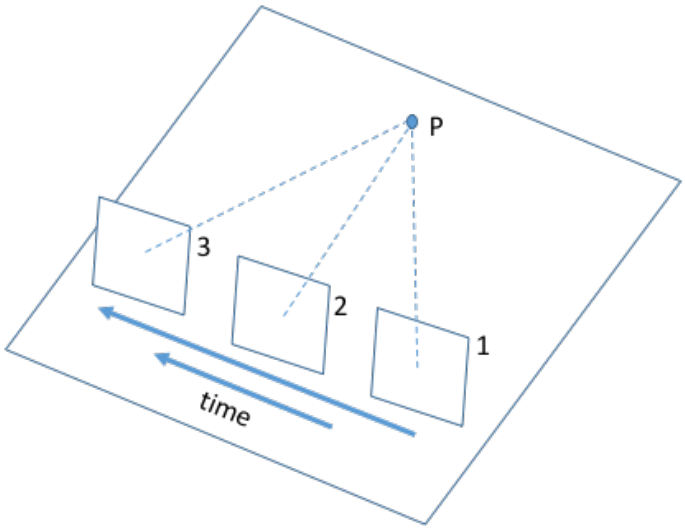

Figure 2.6: Linear capture motion for SFM reconstruction (image adapted from Bolles et al. (1987))

Capture motions for the image set are not limited to linear motions; other common motions include partial or full circumvention of the object of interest. Bolles et al. (1987) noted that SFM 3D reconstruction can be executed in two ways. The first is SFM using knowledge of the camera motion. This technique is usually implemented in applications where the position of the camera is accurately known in the real-world 3D space. The second approach uses matching features in the image set to estimate the camera motion in the real-world (position and orientation of each frame) and is commonly used in automated photogrammetric processes where the manual input of camera motion data is not possible. In this approach, the SFM reconstruction algorithm matches features between images in the image set and forms a sparse point cloud (image tie-points) in the 3D space. Using this sparse point cloud, the algorithm estimates the position and orientation of each image in the 3D space. Once the positions and orientation of the images are defined, the algorithm implements a more rigorous feature matching process and estimates the surface of the object by a dense point cloud. Finally, the surface of the object is reconstructed by applying a mesh over the dense point cloud.

Accurate structure-from-motion (SFM) reconstruction is dependent on the similarity of the bundle of images used as the input data set, enabling easier feature recognition throughout the process (Schenk, 2005). Sufficient overlap (50-80%) between images in the image set is

required. In a paper outlining the photogrammetric process, Schenk (2005) noted that camera calibration also form an integral part of image bundle similarity. Furthermore, he outlined the elements of camera calibration as the coordinates of reference points (thus the real-world distance between them), the calibration of camera focal length and the image quality measures (i.e. resolution).

2.3.3 Health-related applications of close-range mobile phone photogrammetry

The use of mobile phone photogrammetry has been an area of interest to academics since 2008 (Smith and Kokkas, 2012). Due to the increased use of smartphones, Smith and Kokkas (2012) conducted research to compare close-range photogrammetric capabilities of mobile devices (iPad 2 and iPhone 4) and an SLR digital camera. Using these devices, images of a building were captured and a 3D reconstruction of the building was created using PhotoModeler (PhotoModeler, 2012). It was found that although the mobile devices have smaller lenses and sensors, they still have good photogrammetric potential. In comparing the 3D model produced using the SLR camera and the 3D model produced using the iPhone, it was found that there was an average standard deviation of 0.04% in measurements.

Facial prosthesis modelling using monoscopic photogrammetry

Salazar-Gamarra et al. (2016) developed an application and image acquisition technique to 3D model the patient's face as a tool to 3D print facial prosthesis for patients with maxillofacial defects. The images used in the study were captured using a mobile phone (Samsung Galaxy Note 4: www.samsung.com) and were uploaded to an internet server for processing and 3D modelling. To ensure that the 3D model of the patient's head was correctly scaled the images were taken at known distances and angles from key points on the face of the patient and other reference markers such as the floor and backboard. It was ensured that the images were taken with sufficient light on the patient's face so as to avoid the formation of shadows. The images which were used were manually inspected for overexposure or underexposure and the presence of shiny surfaces, this resulted in a selection of 15 2D photos which were used for the construction of a 3D model. This stringent protocol ensured an accurately scaled model of the patient's face which was used to 3D print a facial prosthesis.

Smartphone assessment of diabetic foot ulcers and wounds

Foot ulcers, which can be painful and are prone to infection, are common in individuals with type 2 diabetes. These kind of wounds are evaluated by a clinician on a visual basis without recording quantitative parameters thus making it difficult to assess the healing process of the ulcer (Sudheer, 2016). Sudheer (2016) set out to develop a solution to this problem by developing (a) an imaging protocol for patients to image their foot ulcers using a smartphone and (b) an algorithm to perform real-time analysis of the wound to extract relevant wound data. The images of the ulcers were captured by use of an "image capture box" which is a rig that positions the foot and the smartphone a set distance apart. By means of lighting fixtures on the image capture box, the required illumination conditions are achieved in the captured images. The picture of the ulcer is then segmented using the K-means clustering method to distinguish between the ulcer and surrounding foot. Using the segmented image and the known distance of the camera from the foot, the area of the ulcer is determined and used to track the healing process of the wound.

Similarly, Varma et al. (2016) developed a smartphone application to eliminate the use of contact wound assessments such as the ruler, graph and acetate methods. Contact wound assessments require the clinician to physically touch the wound, resulting in inaccuracies and also causing significant patient discomfort (Varma et al., 2016). The application was developed with the idea of allowing the patient to image their wound at home or at work with their personal smartphone. The user is then required to estimate the distance between the wound and the camera lens (x). The distance x and the known horizontal and vertical camera viewing angle of the smartphone were used to estimate the area of the wound. This method of area calculation was not extensively tested; however, it was concluded that accurate wound area measurement is highly dependent on the accuracy of the estimated distance between the wound and camera lens.

2.3.4 Summary of mobile-based photogrammetry

SFM photogrammetry is a powerful tool for reconstructing objects in the 3D space and deriving real-word information from them. However, the accuracy of information acquired from 3D reconstructions of an object is dependent on various parameters of the images in the image set such as image resolution and image overlap. These parameters are

particularly important for automated 3D reconstruction where feature matching plays an important role in determining image positions and orientations in the 3D space.

Image acquisition is a pivotal component of accurate photogrammetry as it informs many of the factors involved in performing photogrammetry. These include ensuring high image quality for accurate reconstruction and real-world reference data for accurate scaling. Development of a robust image acquisition protocol which accounts for the requirements of the unique problem at hand is pivotal in pursuing the desired outcome of any photogrammetric application.

3. DEVELOPED SYSTEM OVERVIEW

3.1 Design criteria

On consideration of the literature review, the following design criteria for the development of a suitable solution for mobile based TST evaluation were established to support the objectives of the project as outlined in Chapter 1. The solution would:

- Be developed on a platform that is affordable, accessible, and useable in both developed and developing countries – Objective 1 and 2.
- Require minimum data costs with regard to transmission of induration images and patient information – Objective 2.
- Provide a fully automated evaluation procedure – Objectives 3 and 4.
- Produce results with an intra-observer standard deviation of less than 1.1 mm – Objective 4.
- Be developed as a screening tool and should archive the patient information, intermediate results, and final results for review by a trained professional - Objectives 3 and 4.

3.2 System overview

To fulfil the criteria set out above, a smart phone application was developed for patient information collection and image acquisition on the android platform. In addition, an automated method for 3D reconstruction of the induration and extraction of relevant information from the 3D reconstruction was developed. Figure 3.1 summarises the flow of functions of the developed system, which is separated into three modules that represent objectives of the project: 1) image acquisition protocol and application interface development; 2) automated 3D reconstruction; and 3) automated data extraction from the 3D reconstruction and induration measurement algorithm.

The system comprises a two-part process implemented in three modules (Figure 3.1): a mobile unit (module 1; a smartphone) facilitates image acquisition; a remote processing unit (module 2; a personal computer) carries out 3D reconstruction; and measurement and analysis takes place in module 3. The mobile unit is broken down into two major components; the application interface and image acquisition. The application interface facilitates retrieval of relevant patient information and guides the image acquisition

protocol. The mobile unit would be used to input patient-specific information and capture the required images (guided by the developed mobile application interface). These images are sent to the remote processing unit for completion of the evaluation process.

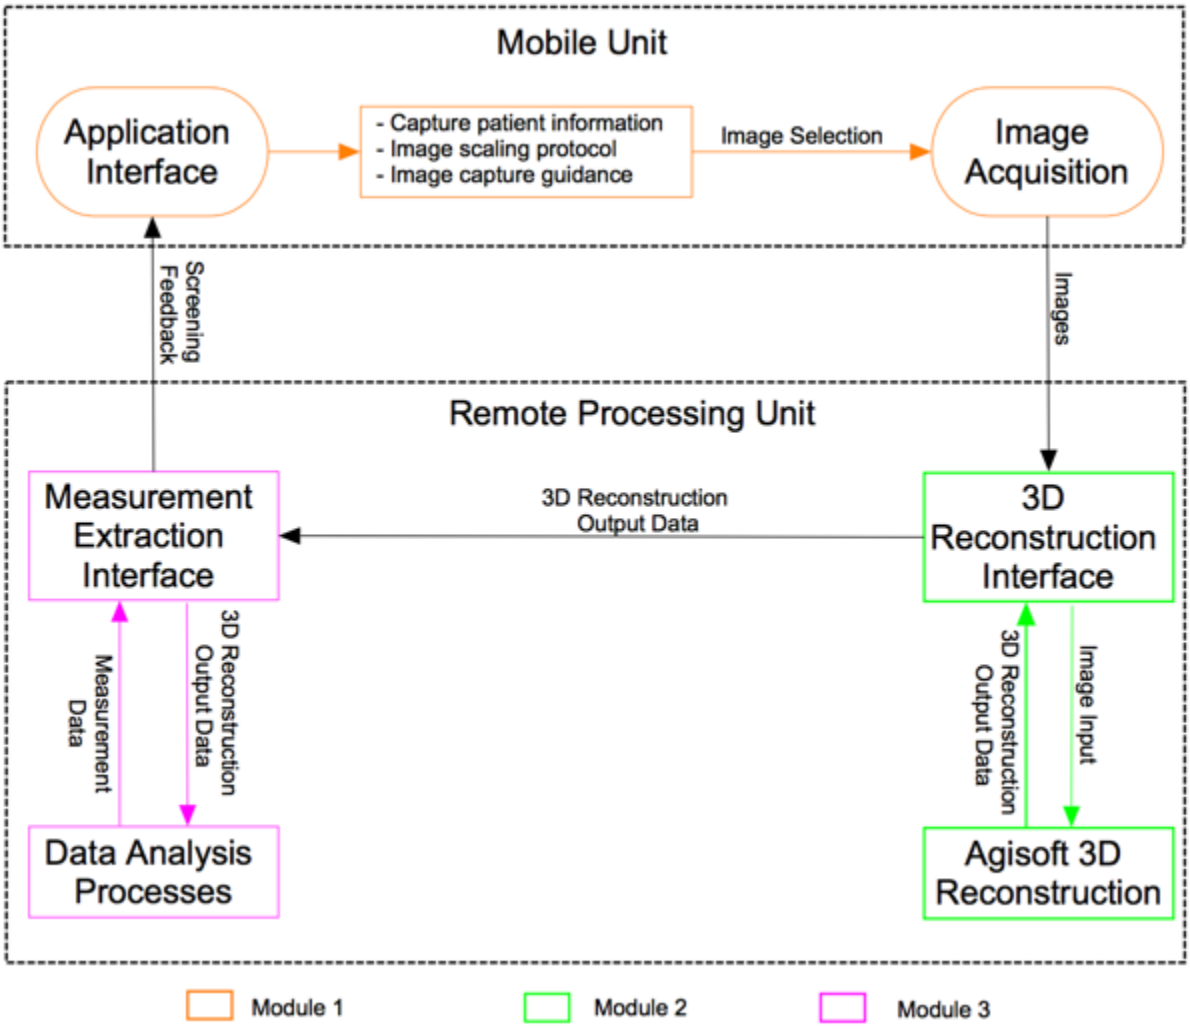

Figure 3.1: Function flow chart of system solution.

The remote processing unit was designed to be, in its most basic form, an internet-enabled computer equipped with Python and Agisoft PhotoScan. The images and information received from the application are used for photogrammetric 3D reconstruction and assessment of the result. The 3D reconstruction module (module 2) was designed as an automated process of photogrammetric reconstruction using Agisoft PhotoScan. Finally, the assessment of results from the 3D reconstruction module (module 3) utilises the 3D reconstruction from module 2 to determine the size of the induration. The results of the process are both stored on the processing unit (for review by a clinician) and sent back to the patient. The system design criteria presented in this chapter influenced the construction of the overall system design. However, these criteria also informed specific design decisions which will be outlined in the chapters to come.

4. IMAGE ACQUISITION: TECHNICAL ASPECTS

This chapter describes the methodology and design of the image acquisition protocol with reference to the technical aspects which were considered. The development of the protocol was informed by factors relating to the clinical evaluation of the TST, the photogrammetric 3D reconstruction of the arm and the evaluation of results from the 3D reconstruction. A unique scaling method was developed for use in this project. This chapter outlines all the above and each subsection concludes with a description of the integration of the respective component into the final image capture protocol.

4.1 Image acquisition considerations for clinical evaluation of the TST and photogrammetric reconstruction

4.1.1 Considerations for clinical evaluation of TST

Evaluation of the TST has specific features which inform how the induration is measured. Since the size of the induration is defined to be the width of the induration measured perpendicular to the long axis of the arm, it follows that two features are of significance: the long axis of the arm and the edges of the induration. Figure 4.1 shows an arm with an induration and illustrates the long axis of the arm.

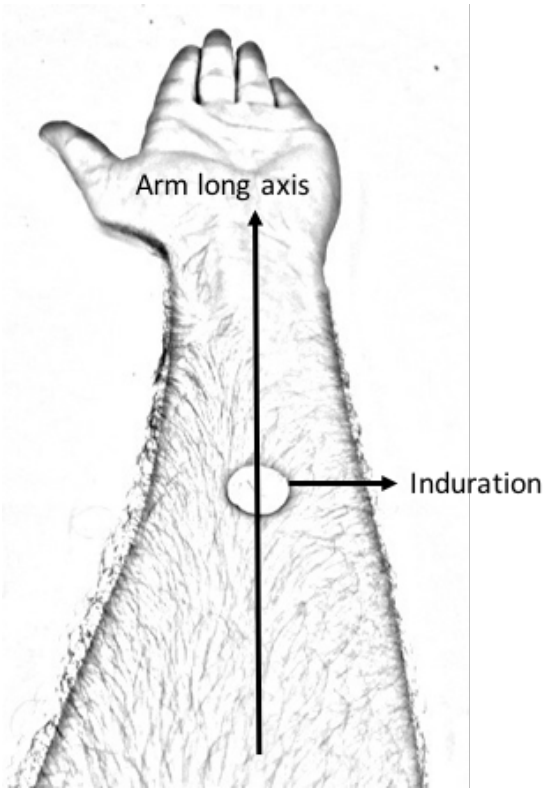

Figure 4.1: Long axis of forearm showing induration.

The two features illustrated in Figure 4.1 need to be clearly identifiable in the captured images. With reference to the long axis of the arm, it is important to note that when the size of the induration is evaluated by the clinician, the direction of this axis is guided by the intuition of the clinician. However, for this project, in the absence of a clinician, the direction of this axis would be guided by the intuition of the person capturing the images. Figure 4.2 shows two images of an arm with an induration.

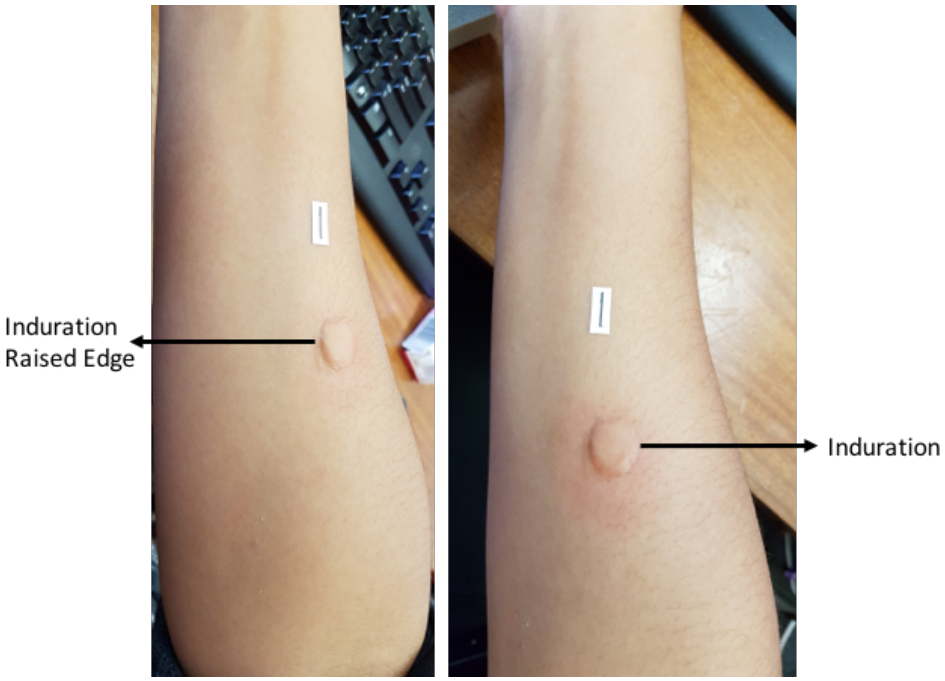

Figure 4.2: Image of top view of an arm with induration (right) and left-tilted image of induration (left).

In Figure 4.2, the right image depicts the top view of the arm (perpendicular to the surface of the forearm) while the left image depicts an image of the same induration slightly tilted to the left of the induration. The start of the induration (points from which diameter is measured) is defined as the point where the induration rises off the surface of the skin of the forearm. For the evaluation of the TST, the start of the induration perpendicular to the long axis of the arm is of particular interest. As can be seen in the top view image in Figure 4.2, the raised edges of the induration are difficult to identify. This is because very little depth information is available in this image. However, the tilted image (Figure 4.2 – left) shows the raised section of the induration which would contain the point perpendicular to the long axis of the forearm. To acquire sufficient depth information to evaluate the size of the induration in accordance with the methods of evaluating the TST, it is desired that as

much information about the raised sides of the induration perpendicular to the long axis of the arm be captured.

4.1.2 Considerations for photogrammetric reconstruction and result evaluation

4.1.2.1 Image capture environment and lighting conditions

As stated in Gamadia et al. (2007), low light conditions have a significant influence on the image quality and noise present in the captured images, and affect the quality of the 3D reconstructed model. Imaging environments with too much light also cause problems, such as overexposure, which may result in a loss of detail in the image. For the purpose of repeatability and consistency of image quality, suitable lighting conditions should be specified in the imaging protocol.

4.1.2.2 Photogrammetric 3D reconstruction

For best results in any structure-from-motion (SFM) photogrammetric process, high image resolution is desired. This is because the basis of the process involves matching various properties such as colour and texture of pixel sets between images in the image set. High-resolution images capture these properties more clearly resulting in more tie-points (matching features) between images and thus a better 3D reconstruction of the object of interest. Furthermore, successful photogrammetric 3D reconstruction requires sufficient overlap between images in the image set. It is common practice that the images in the image set have 50-80% overlap (Walford, 2007).

4.1.2.3 Result evaluation process and scaling

In an automated evaluation process, it is imperative that features relevant to the evaluation process are scaled and identifiable by the implemented automated evaluation algorithm. As shown in Bargal et al. (2015), Maier et al. (2015) and Wallis et al. (2016), this can be achieved by specifying positioning and alignment constraints in the image acquisition protocol.

4.1.3 Constraints on photogrammetric reconstruction and result evaluation

4.1.3.1 Image capture environment lighting conditions

Considering that the developed mHealth solution is to be used in various environments, it is difficult to ensure adequate lighting conditions. Due to this restriction, adequate lighting conditions can be achieved through two avenues: using some light source which would be

accessible at all locations or specifying some measure of good ambient light conditions in the environment in which the images are captured. Given that most smartphones are equipped with a flash for capturing images, the light from the flash could be used to mimic standard lighting conditions for the image capture protocol.

4.1.3.2 Photogrammetric 3D reconstruction

Agisoft PhotoScan (APS) was used for photogrammetric reconstruction in this project. The APS user manual suggests the use of cameras with at least 5 megapixel resolution, which corresponds to the camera resolution of the HTC One smartphone used in this project. Furthermore, when using cameras of resolution close to the lower bound of 5 megapixels, large image sets are desired for accurate reconstruction. This is because larger image sets provide more information about the object of interest. This results in a greater number of tie-points and larger 3D dense point clouds (point cloud estimation of the reconstructed surface), which in turn increases the accuracy of reconstructed features. However, with a larger image set comes increased data costs associated with the transfer of a larger number of images.

Alongside the number of images in the image set, it is also important to ensure that there is a 50-80% overlap between images. To fulfil this constraint, APS's built-in 'image quality' measurement can be used. An image quality measurement of greater than 0.5 indicates sufficient overlap and clarity between successive aligned images.

Finally, limiting background objects when capturing images is desired. Although this could be achieved by cropping the images post-acquisition, a better alternative is capturing photographs of sufficient overlap (50-80%) close to the object of interest. In such a case, the photogrammetric process focusses more on reconstructing the object of interest (as opposed to other objects in the image field of view) thus reducing the processing time. In addition, capturing photographs close to the object of interest is advantageous when using low-resolution (5 megapixels or less) cameras. This is because these photographs reveal more detail used in the SFM point matching process than photographs taken further away. Finally, since automated photogrammetric processes utilise surface feature detection to match tie-points across the image set, it is important that the features of interest in the image set are clear and in focus.

4.1.3.3 Result evaluation and scaling

In order to meet the design requirement for a fully automated process for the TST evaluation, it is imperative that the location of the induration and direction of the long axis of the arm are specified and identifiable by the evaluation algorithm. Furthermore, scaling of the image set to real-world scale is a key contributor to a successful photogrammetric reconstruction and extraction of real-world measurements. It is thus also important to note that the image acquisition protocol should be designed to accommodate a scaling method for the image set.

4.1.4 Developed image capture protocol

The developed application implements the image capture protocol and guides the user in capturing 7 images of the induration guided by an orientation panel and guide boxes on the image capture screen. In order to determine the appropriate number of images for the evaluation process, an accuracy experiment was conducted (presented in Chapter 8) which compared the use of 3, 5 and 7 images in the reconstruction process. When evaluating the additional data cost (for mobile internet transfer of an increased number of images) and the required measurement accuracy of the induration, it was found that a 7-image protocol was most appropriate for use in this project. Figure 4.3 illustrates the general orientation and position of the 7 images and the image capture screen with guidance components.

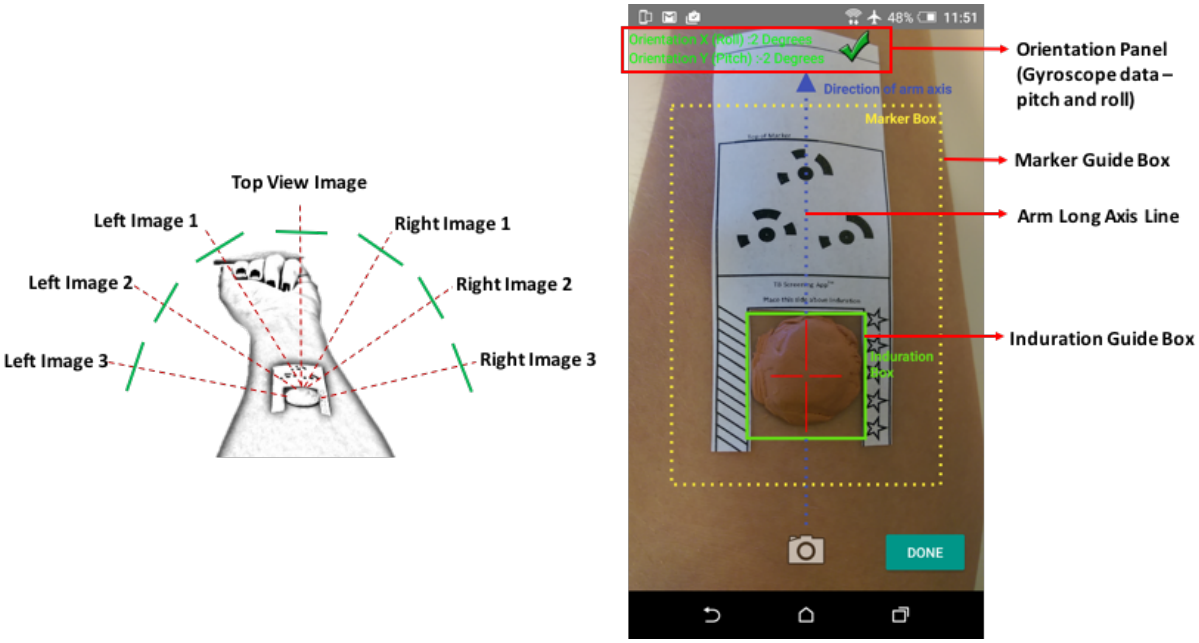

Figure 4.3: Illustration of position and orientation of 7-image imaging protocol and a screenshot of application image capture screen.

Four guide components shown in the screenshot of the application image capture screen (Figure 4.3 – right) were developed to guide the image capture process. The following is a general overview of the function of each of the components:

- Orientation panel: Shows the gyroscope pitch (tilt forward/backwards relative to the horizontal) and roll (tilt left and right relative to the horizontal) data in degrees.
- Marker guide box: The scaling marker is contained in this guide box during the image capture process to ensure that the entire marker is included in the captured images.
- Arm long axis line: Denotes the orientation of the long axis of the arm.
- Induration guide box: The induration is contained in this guide box during the image capture process. This bounding box guides the user to ensure that the induration is completely visible and in the correct location (in the image) during the image capture process.

4.1.4.1 Image capture specifications

Built-in features of the developed application were designed to ensure that the images were correctly captured in a manner best suited for the reconstruction process. The autofocus feature was used to keep each image in focus with the autofocus weighting skewed toward the pixels in the induration guide box. The capture process was designed such that the camera flash was set to torch mode (flash is on during the time the user is orientating the device), which allows the autofocus feature to account for the light from the flash when focusing. This ensured that the induration, and skin which surrounds it, were clear and in focus.

Each of the 7 images required for completion of the image capture protocol has unique specifications which satisfy the requirements for 3D reconstruction and evaluation of the result. Figure 4.4 shows the image capture specifications for the images captured using the developed application.

Figure 4.4: Image capture specifications for the 7 images captured using the mobile application.

Figure 4.5 below shows examples of each of the 7 images captured in accordance with the specifications outlined in Figure 4.4.

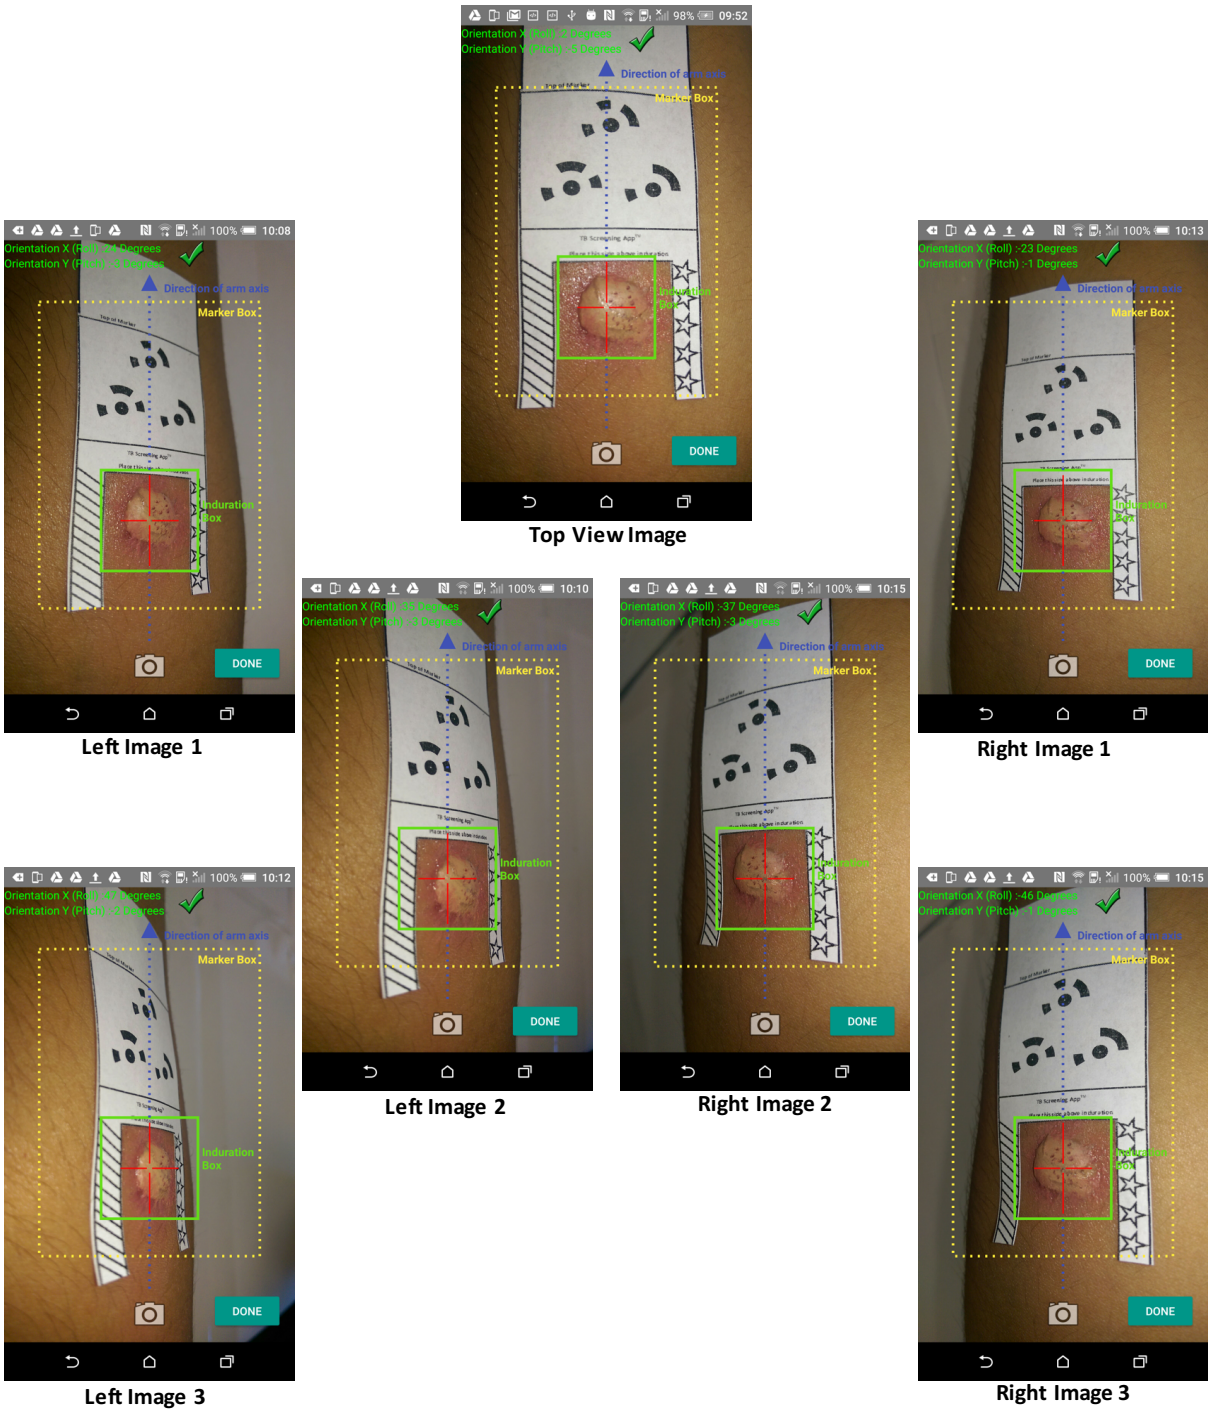

Figure 4.5: Example images of 7 images captured using the mobile application.

The specified angles for each image allow for a total coverage of 85° to 100° of the forearm and ensure the required 50-80% overlap between successive images. Furthermore, the designed image specifications encourage the user to capture images as close as possible to

the surface of the arm. As seen from Figure 4.4, the specification for the top view image is more detailed than for the other six images. The reason for this is that this image is used as a reference for the location of the induration and the long axis of the arm in the 3D reconstruction and evaluation processes.

4.2 Scaling methods

A pivotal aspect of photogrammetry is scaling the images used in the photogrammetric process, and this needs to be integrated into the image capture protocol. Accurate scaling of an image set allows for accurate measurements to be made in the 3D space. This section will begin by summarising the various scaling methods which were considered. This is followed by a detailed description of the design and development of the scaling sticker which was used in this project.

4.2.1 Video scaling method

When using a smartphone to capture images for photogrammetric reconstruction, one can make use of sensors embedded in the mobile device to solve for the scale of objects in the images. Ham et al. (2015) developed a novel solution for scaling objects recorded on video by a smartphone (iPhone 6) using Inertial Measurement Unit data (IMU) from the phone as the video is being recorded. The IMU of any device comprises of a set of gyroscopes and accelerometers which measure the linear and angular acceleration of the smartphone at different points in time. Aligning this data with single frames from mobile phone video footage, the authors were able to solve for the scale of objects in the video by monitoring the change in the size of the various objects in the frames and aligning this to the change in displacement of the mobile device (derived by double-integrating the linear and angular accelerations of the mobile phone).

Ham et al. (2015) implemented this scaling method on various objects using video and IMU data from an iPhone 6. It was found that a 15-second-long video at 240 frames per second (fps) resulted in an output scale with less than 2% error. Upon attempting to implement the algorithm (for use in this project) using IMU data and standard video from the HTC One M8 at 120 fps, it was found that the time stamps of the IMU data did not align with the frames of the video input. This drawback of the data input was cited by Ham et al. and was mitigated (in a pre-processing step) in their study by only selecting frames from the video

which had time stamps that matched the IMU data. However, the authors also indicated that the high frame rate of the video from the iPhone 6 made it possible to achieve a low scaling error and that video frame rate plays a role in the scaling result.

The video scaling method aligns with the two-phase system outlined in Chapter 3 as the video and IMU data can be recorded on the device and sent to the central processing unit to be scaled. However, the data costs for transferring the video and IMU data to the central processing unit reduces the practicality of this approach. A 15-second-long video recorded on the HTC One and IMU data for a similar time period is typically 20 to 30 megabytes (MB) and 2 MB in size respectively. This would be costly to transfer over a mobile internet connection. Although the advantages of using this scaling method are significant, it was rejected for the flowing reasons:

1. Limitation on the specifications of the mobile phone that can be used: The accuracy of the scaling process decreases with decreasing frame rate of the recorded video. It was therefore determined that the scaling method would not be compatible with low-cost smartphones used in resource-limited settings.
2. High data transfer costs: The cost of transferring sizable video and IMU data to the central processing unit would be high. This high data transfer cost was undesirable as this burden would be placed on the end user of the mHealth application and thus limit the use of the application in resource-limited settings.

4.2.2 Reference length scaling method

The use of a reference length to scale objects in an image is the most commonly used method of scaling for photogrammetric processes. Various projects (Chen et al., 2013, Salazar-Gamarra et al., 2016, Varma et al., 2016) used reference length scaling in photogrammetry to scale image sets captured on a mobile phone, by inclusion of an object of known length in the captured images. Using this known length in the image, the ratio of the length of two objects in the picture can be compared with the ratio of the length of the same two objects in the real world and thus any object in the image can be measured. The advantages of this scaling process lie in its simplicity, accuracy and low implementation cost. However, all reviewed mobile phone photogrammetric applications of this method (Chen et

al., 2013, Salazar-Gamarra et al., 2016, Varma et al., 2016) involve manually identifying the object of known size in each image. This would be incompatible with automated scaling; however, this problem could be solved by replacing the object of known distance with a set of coded-markers with a known distance between them. This would allow for automated identification of the known length and thus automation of the scaling process.

4.2.3 Developed scaling solution: scaling sticker

For use in this project a unique scaling sticker, which implements reference length scaling, was designed. This subsection will outline the design considerations which influenced the development of the scaling sticker followed by a presentation of the developed scaling sticker.

4.2.3.1 Scaling sticker design considerations

The following design criteria were specified for the development of the scaling sticker used in this project. The scaling sticker should:

- Accommodate the range of expected induration sizes.
- Implement a coded marker base reference length scaling method.
- Be intuitive to use.
- Be easily distributable.

Although the primary purpose of the scaling sticker is to scale the image set, it was found that that the sticker could be designed to further complement the 3D reconstruction and results evaluation processes. The first of these additional functions of the scaling stickers was the specification of the arm long axis orientation. This scaling sticker was designed to be used as a second arm long axis (in addition to the capture screen arm axis) used to measure the induration. The second additional feature was aimed at improving the result of the photogrammetric 3D reconstruction. As shown by Salazar-Gamarra et al. (2016), adding markers or distinct features in the capture space could improve the photogrammetric 3D reconstruction process. This is because the photogrammetric reconstruction software recognises more distinct features in the region of interest, thus resulting in a better reconstruction of the object. While coded markers are commonly used for achieving this effect, any distinct shapes or features added to the scene can be used.

4.2.3.2 Developed scaling sticker

Figure 4.6 shows two images of the developed scaling sticker, with and without labels of its various components and dimensions.

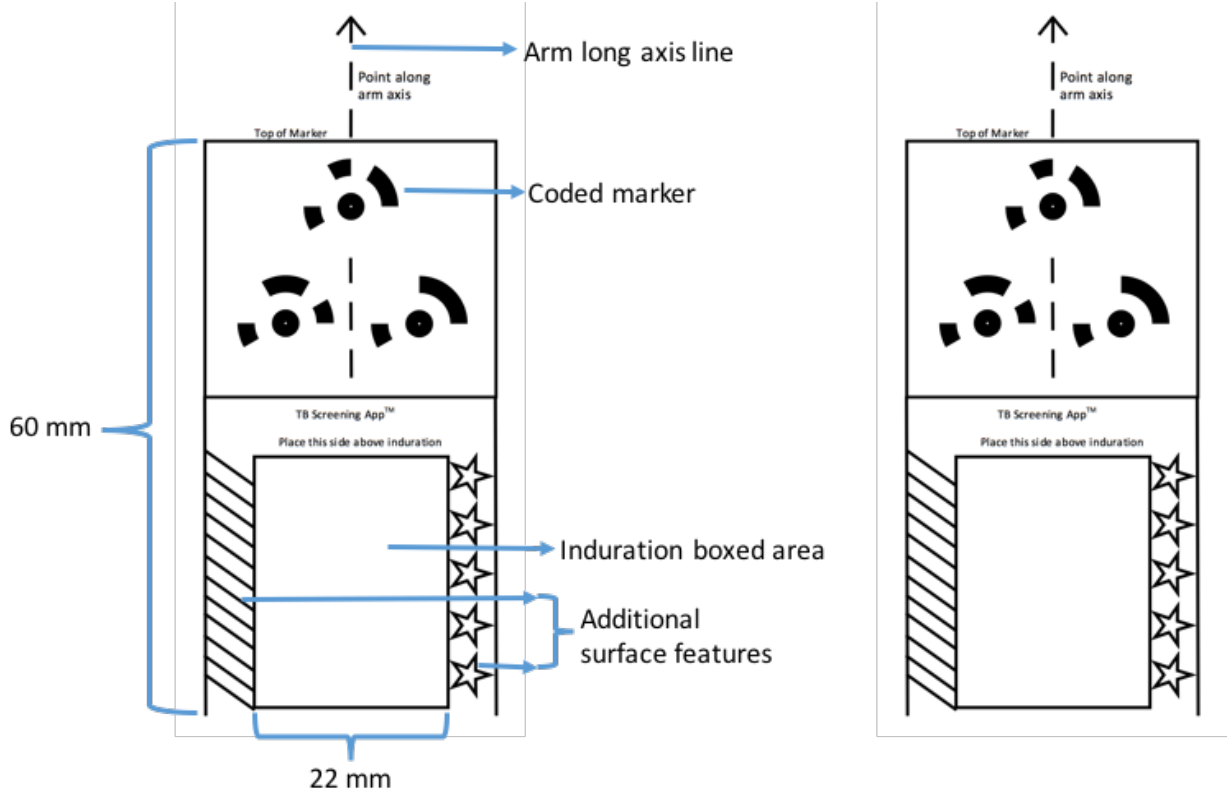

Figure 4.6: Scaling sticker with labelled features (left) and without labelled features (right).

The scaling sticker shown in Figure 4.6 was designed to be stuck on the forearm of the person being evaluated using the mHealth application. The induration boxed area is 22 millimetres in width to accommodate the upper bound of induration sizes. The marker is made up of four major components: arm long axis line, coded-markers, induration boxed area and additional surface features.

The subsections below outline the function of each of these components.

4.2.3.3 Coded-markers and arm axis integration

The coded-markers used in the sticker design were obtained from the APS library and can be detected by the program automatically. Figure 4.7 shows the coded-markers on the scaling sticker.

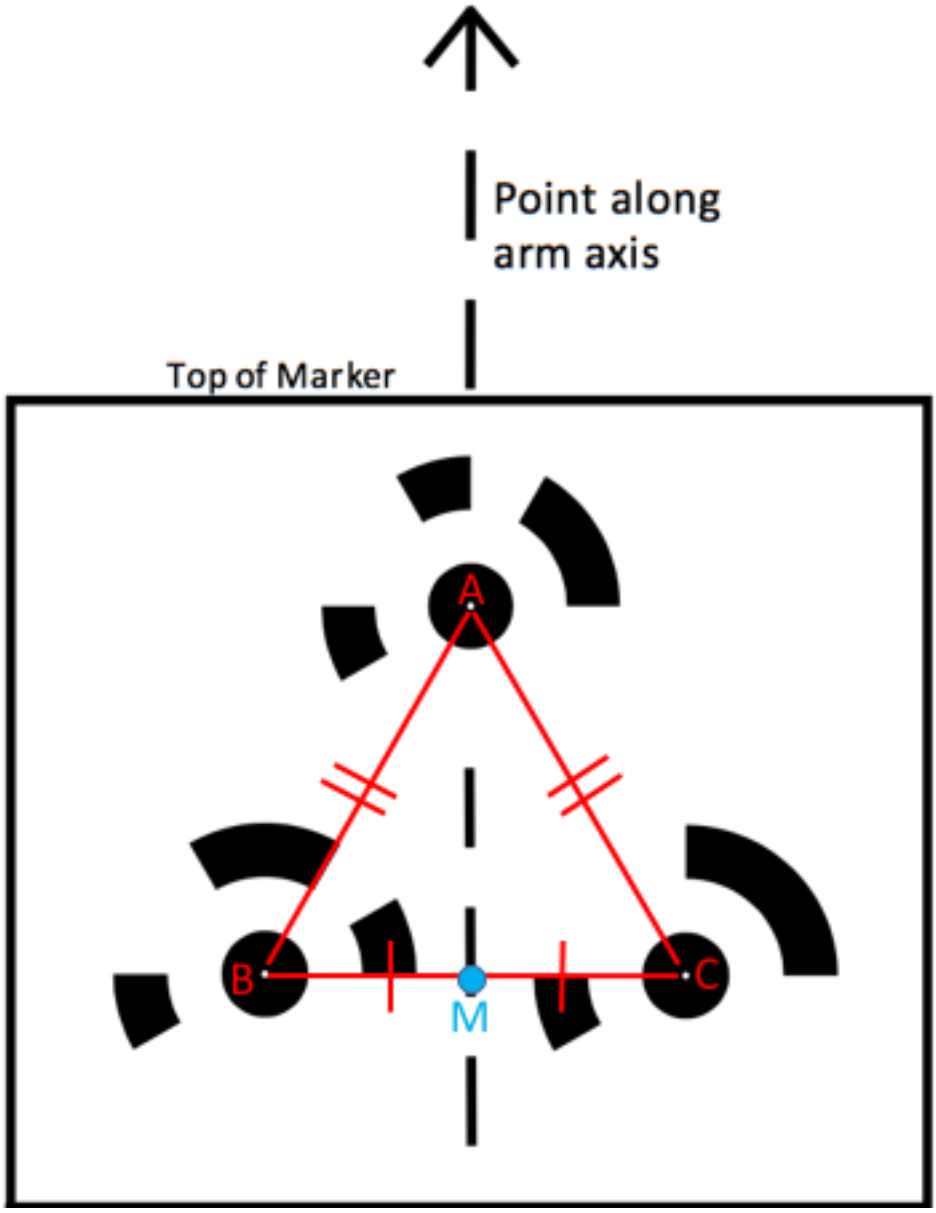

Figure 4.7: Scaling sticker coded-markers in an equilateral triangle formation.

Each of the markers is 3 mm in diameter and they are arranged in the formation of an equilateral triangle with sides of length 14.5 mm. As seen in Figure 4.7, the coded-markers are labelled A, B and C respectively with point M representing the midpoint between markers B and C. The imaginary line between coded marker A and point M forms the arm long axis line which is represented by the dashed line extending out of the box containing the coded-markers. Using the aforementioned points (A, B, C and M), the evaluation algorithm identifies the long axis of the arm from the orientation of the scaling sticker in the final phase of the automated evaluation process. The orientation of the long axis of the arm is guided by the intuition of the patient. In this case, it can be determined by the perceived

long axis of the arm when applying the scaling sticker and is referred to as the scaling sticker arm long axis.

4.2.3.4 Additional surface features

Figure 4.8 shows the lower half of the scaling sticker, which contains additional surface features surrounding the induration boxed area.

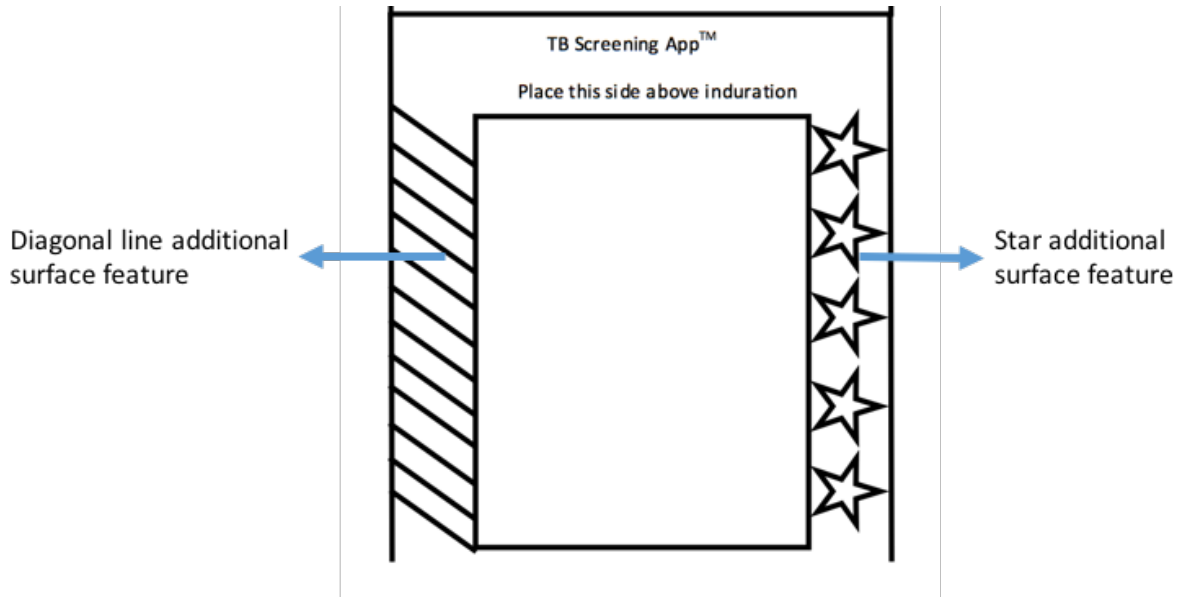

Figure 4.8: Lower half of scaling sticker showing additional surface features of the scaling sticker.

One of the key factors in evaluating the size of the induration is identifying the margins of the induration (where the induration starts and ends). As shown in Figure 4.8, diagonal line and star surface features flank the left and right side of the induration boxed area, respectively. The shapes (lines and stars) of the chosen additional surface features were arbitrary, but emphasis was placed on the location of these features. Since the features are located along the flanks of the induration, they will mark the skin surrounding the induration. This will result in more tie-points across the image set in this area and thus a better 3D reconstruction of the surface surrounding the induration.

5. IMAGE ACQUISITION: USER AND ENVIRONMENTAL CONSIDERATIONS

The greatest variability is introduced into the processing pipeline during the phase of image acquisition. This variability could present as blurry images (due to instability at the point of capture), insufficient lighting, and misinterpretation of the image capture protocol. The purpose of developing a structured image acquisition protocol is to minimise such variability through making considerations for human-centred design. This chapter outlines the image acquisition design considerations, including an overview of the image acquisition environment specifications.

5.1 Smartphone positioning during image capture

5.1.1 Smartphone orientation

Physically capturing images is a source of possible unwanted instability which could result in blurry images. Sudheer (2016) mitigated the human factor in capturing images for an mHealth application by use of a positioning rig for the imaging device. This approach allows for consistent stability and ease of use. However, a major drawback of this approach is that the required rig needs to be available at all locations where the mHealth application is used. In this project, we aimed to impose minimal equipment, other than the mobile phone, for image acquisition. Given this constraint, the image capture protocol needed to ensure stability for images captured in free-form (without the use of a rig). Users imaging their own induration on their arm will only have one hand available to operate the smartphone. Figure 5.1 shows a common method of capturing images with one hand using an Android device.

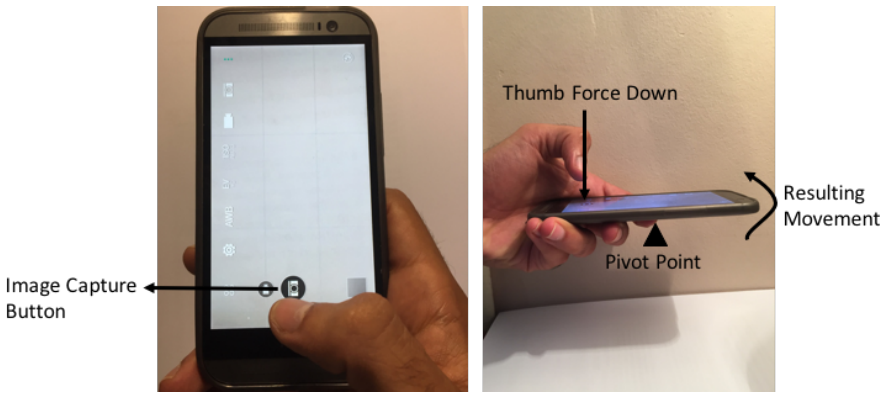

Figure 5.1: Top view (left) and side view (right) of a common method of single-hand image capture using a smartphone.

As shown in the top view (Figure 5.1 – left), the image capture button is situated at the bottom of the touch screen of the smartphone with the other four fingers wrapped around the lower half of the back of the device. This method is, however, unstable. The side view (Figure 5.1 – right) shows how the action of capturing an image using this grip could cause the top of the phone to jerk upward. Since only the lower portion of the back of the phone is supported by the user's fingers, a pivot point is created about the index finger of the operator. As the operator pushes down on the capture button, the phone slightly rotates around this pivot point causing a resultant movement at the far end of the phone (where the camera lens is situated). This could result in a blurry image due to instability at the point of capture. It is thus important to incorporate a method of mitigating such instability in the image capture protocol.

5.1.2 Suggested method of holding device during image capture

Figure 5.2 shows the recommended method of holding the mobile device, to ensure greater stability.

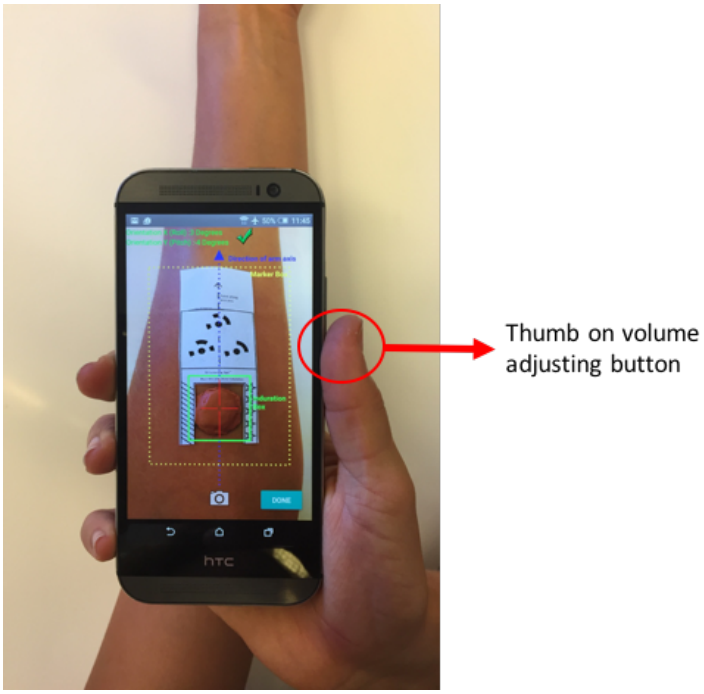

Figure 5.2: Suggested method of holding the mobile device while capturing images using the developed mobile application.

Although the user has the option of capturing the image with the camera icon at the bottom of the screen, the HTC One volume adjustment buttons are configurable within the developed application to be used instead for increased stability at the point of capture. As

shown in Figure 5.2, the device is held with the user's palm behind the mobile phone and thumb on the volume adjustment button (which is used to capture the images). Holding the mobile device in the suggested manner is more stable because the force exerted by the thumb to capture the image is directed across the face of the device. Since the device is fully constrained by the user's fingers on the opposite side of the phone, there is minimal movement of the phone when capturing images.

5.2 Scaling marker - operator considerations

Figure 5.3 below shows two pictures illustrating how the scaling marker sticker is applied to an arm with a simulated induration.

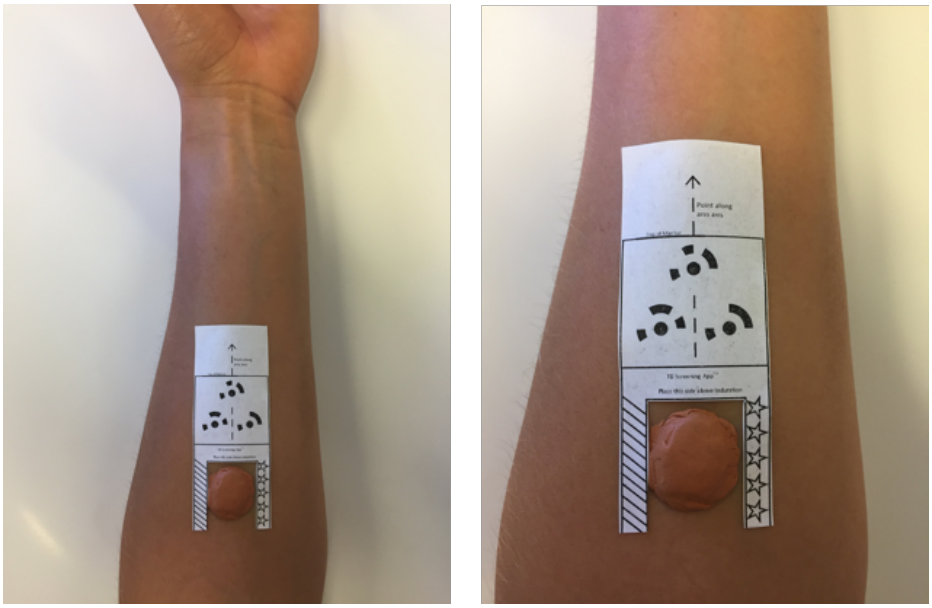

Figure 5.3: Illustration of scaling sticker application on arm with simulated induration.

With instructions from the mHealth application, the scaling sticker is designed to be applied to the skin with the induration contained in the induration boxed area (Figure 5.3). Furthermore, the sticker is to be orientated so that the arm long axis line on the sticker points along the long axis of the arm as best as possible from the operator's perspective. The scaling sticker is printed on sticker paper using a conventional printer (printable sheet shown in Appendix A). This ensures that the design can be distributed and printed at the location where it is required. This low-cost and easily distributable characteristic of the scaling method facilitates wide use of the application in resource-limited settings. In practice, this scaling sticker can be distributed through two avenues. The first of these, pertaining to persons with their own mobile phone, is by suppling patients with a scaling

marker sticker upon administering the TST. The patient is required to apply the scaling sticker to the arm before imaging the induration using the mHealth application. The second avenue, pertaining to persons without a mobile phone, is by supplying the scaling marker to the patients at a central location which is equipped with a mobile phone and the mHealth application. The patients can visit this central location for evaluation of the result of the TST.

5.3 Collection of patient-specific information

Alongside guiding the image acquisition process, the mobile application is also designed for collection of patient information, which requires the following: personal patient information (for identification purposes), the status of patient's health condition and the patient's level of exposure to TB. This information may be used to determine the induration size range which should be interpreted as a positive LTBI result, as indicated in Table 2.1 in Chapter 2.

5.3.1 Welcome page and LTBI questionnaire in developed application

Figure 5.4 below shows the first two screens that the user encounters when opening the application on a mobile phone.

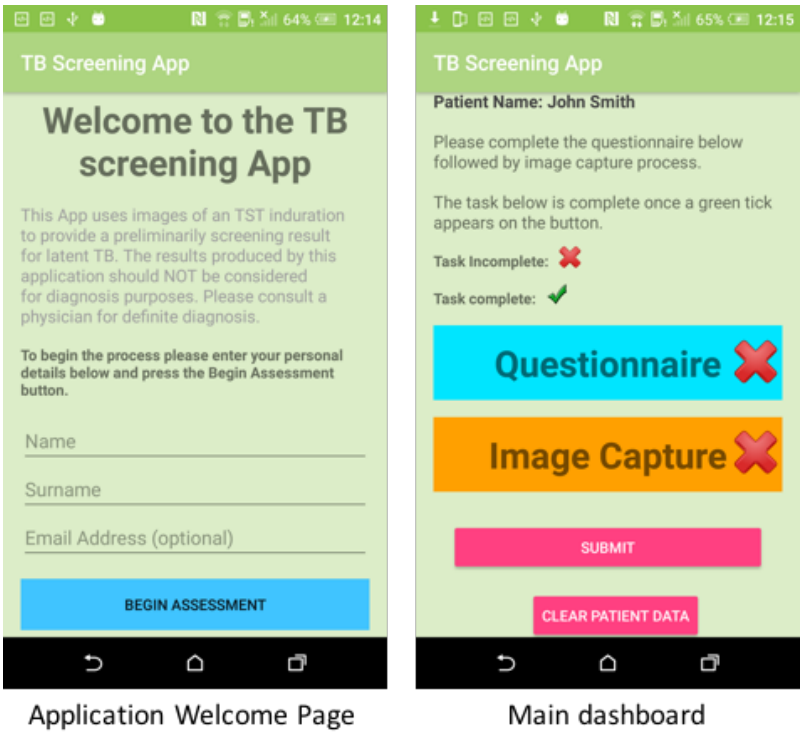

Figure 5.4: Application welcome page and main dashboard.

Upon opening the application, the user is directed to a welcome page where their name, surname, and email address (optional) are recorded (Figure 5.4 – left). In a real-world use of the application, this page could be amended to encompass more patient information which

would be used to identify the patient so he/she may be contacted with the result of the screening process. On the second page of the application (Figure 5.4 – right), the user is prompted to complete two tasks: a questionnaire and the image capture process. The red crosses beside the "Questionnaire" and "Image Capture" buttons indicate that the user has yet to complete those tasks. The questionnaire mimics the questions shown in Table 2.1 in Chapter 2 which were used to determine the induration size range which should be interpreted as a positive LTBI result for the specific patient. Figure 5.5 shows the flow of completing the questionnaire and returning to the Main Dashboard page.

Figure 5.5: Application screenshots illustrating flow of completing the questionnaire.

As shown in Figure 5.5, the user is prompted to complete the questionnaire by selecting 'yes', 'no' or 'unsure' for each of the ten questions. Once complete, the user is directed back to the main dashboard and the responses to the questions in Table 2.1 along with the patient's personal information are recorded in a text file by the application.

5.4 Image capture – operational external specifications

Operational external specifications for the image capture protocol are important to minimise the influence of environmental factors during image acquisition. The goal of any image capture process (for object identification) is to have a clear representation of the object of interest. In the TST application it is thus important that the arm being imaged does not blend into the background of the image. This will ensure that the surface of interest (the arm) is clearly identifiable. Furthermore, since the image acquisition process is guided by

the orientation panel (gyroscope readings), the protocol's operational external specifications must complement the use of the orientation panel.

The developed image capture operational external specifications are presented below together with the rationale for each specification:

- The arm with the induration should be imaged on a flat horizontal surface: Imaging the arm in this position assists in producing stable images as the arm will be supported by the flat surface. Furthermore, the gyroscope displays readings relative to the horizontal (the plane perpendicular to gravitational pull), thus imaging the arm in this position results in consistent and intuitive gyroscope pitch and roll readings.
- The background of the arm being imaged should contrast the skin of the patient: A contrast between the arm and the background will allow for the arm to be easily identifiable in the photogrammetric 3D reconstruction process. This can be achieved by placing a paper, with contrasting colour to the skin tone of the patient's arm, on the flat surface used for imaging the arm.
- The images should be captured in a well-lit room: The exact luminance required is not specified as this could vary from place to place. The camera flash of the mobile device will be used as the primary source of light for the images. A well-lit setting for capturing images, however, will minimise shadows in the images.

5.5 Image capture implementation using developed mobile application

The image acquisition protocol was implemented as part of a mobile application developed on the Android platform, developed with consideration to the ease of implementation. Considering that this mHealth application may be used in resource-limited areas, where the use of mobile phones to capture images may not be common, it was considered important to construct the image capture protocol so that it was intuitive and easy to implement by an untrained operator. This was achieved by integrating clear image capture instructions into the application and through the use of intuitive tools and cues during the image capture process. This section will detail the image acquisition protocol and how it was integrated into the functions of the mobile application.

5.5.1 Image capture

Image capture is initiated by pressing the “Image Capture” button on the Main Dashboard page. The operator is directed to an instruction page, consisting of text and example images which outline four main aspects of the image capture process: how to hold the phone, a general overview of the capturing process, guidance components of the image capture screen and instructions on how to apply the scaling sticker. Once these instructions have been read, the operator is redirected to the image capture dashboard. Figure 5.6 shows a screenshot of the image capture dashboard and the process flow of capturing the top view image using the application.

Figure 5.6: Process flow of capturing top view image.

As shown in Figure 5.6, once the “top view image” button (on the image capture dashboard) is selected the operator is directed to the top view image information page. This page contains text instructions for capturing the top view image accompanied by an example

image of an ideal top view image. Following this instruction page, the operator is prompted to capture the image through the image capture screen. The captured image is then displayed to the operator in the image verification screen where the operator can choose to retake the image (if the image is not within stipulated specifications) or save and proceed with the image capture process. If the image is within the stipulated specifications, the operator is directed back to the image capture dashboard where a green tick signifies that the image has been captured. The operator then proceeds to follow the same process for the six remaining images.

5.5.1.1 Image capture orientation guidance

A critical aspect of the image capture protocol is the orientation of each image guided by the orientation panel. Figure 5.7 below shows an example of the application guidance components which assist the operator in attaining values for pitch and roll within the specified range for the particular image being captured.

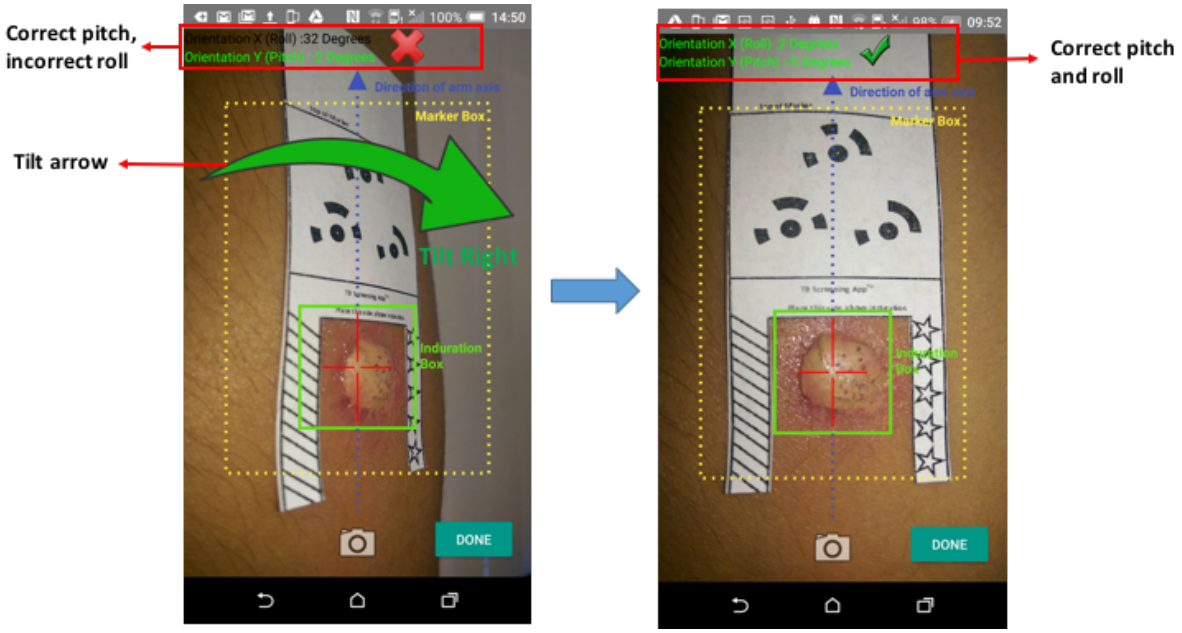

Figure 5.7: Orientation correction features of the developed mobile application.

The specified pitch and roll ranges for the image being captured in Figure 5.7 are both -5° to 5° (top view image). As shown in Figure 5.7 (left) the roll angle is outside the specified range (indicated by black text for the roll reading and a red cross beside the orientation panel). The image cannot be captured at this stage as the application is designed to not allow the operator to capture the images until the device is orientated in the correct pitch and roll ranges for the image being captured. At this point, the operator is directed (by the tilt arrow

shown in Figure 5.7 - left) to tilt the device right toward the correct roll range. When the pitch and roll are within the specified ranges for the image being captured, the text in the orientation panel turns green and a tick appears beside this text – this indicates to the operator that the device is correctly orientated and image may be captured.

5.5.1.2 Application output package

Once all 7 images are captured, a package of data consisting of the following is sent to an online database:

1. 7 captured images in JPEG format.
2. A text file containing pitch and roll values for each image.
3. A text file containing responses from the questionnaire.
4. A text file containing patient information and contact details.

The total size of the package of data varies between 8 and 10 MB.

6. 3D AGISOFT RECONSTRUCTION AND DATABASE DESIGN

This chapter describes the Agisoft PhotoScan (APS) 3D reconstruction and the database design for managing input data and the results of the TST evaluation process. The chapter begins with a description of photogrammetric 3D reconstruction with an emphasis on automated reconstruction in APS including parameters and settings used in the reconstruction process. Following this, the outputs of the reconstruction, which are used in the automated evaluation process, are described. Finally, an outline of the automated construction of a database for each package of information received from a patient through the developed application is presented.

6.1 3D reconstruction using structure-from-motion

Image based SFM programs and algorithms are widely used for a variety of 3D reconstruction applications (Ballabeni et al., 2015). The basic premise of these programs is that a set of images of an object can be used to 3D reconstruct the object if there is enough overlap and depth information about the object in the set of images. Some of the most commonly used programs are Visual SFM (Wu, 2011), Apero MicMac (Deseilligny and Clery, 2011), Eos PhotoModeler (EOS Systems Inc.: www.photomodeler.com) and APS (Ballabeni et al., 2015). These programs all have their basic reconstruction methodology rooted in the process of SFM which uses multiple images at various positions in space to align tie-points (corresponding points in images) and reconstruct objects in the images (Ballabeni et al., 2015). Once images have been aligned, these programs use the identified tie-points and positions of the images to construct a dense 3D point cloud. Finally, using this point cloud an estimation of the surface between points can be approximated and thus, a 3D mesh formed as a surface over these points.

6.2 Photogrammetric SFM reconstruction using APS

APS has been used in the healthcare sector and other industries where accurate 3D reconstruction is necessary (Salazar-Gamarra et al., 2016, Surový et al., 2016). Upon comparing the accuracy of 3D reconstructions from the abovementioned SFM 3D reconstruction programs to laser scanned 3D reconstructed objects, Ballabeni et al. (2015) found that APS consistently had the best performance in terms of image alignment and point cloud generation (object surface approximation). In addition, APS also has the

functionality of interfacing with Python scripts, thus it is possible to automate the 3D reconstruction process for a batch of image sets and automate the extraction information from the resulting model.

The quality of the 3D reconstruction is judged by how accurately the 3D reconstruction represents the surface and features of the reconstructed object. In the general sense, two parameters determine the quality of a 3D reconstruction: the number of tie-points and the number of points in the dense point cloud. Using the tie-points and dense cloud, APS constructs a mesh which estimates the surface of the object of interest. The 3D reconstruction is produced by importing the image set into the APS environment (step 1), initiating the processes of aligning the photos (matching points between the photos) (step 2), generating a dense point cloud (step 3) and building a mesh using the dense cloud points (step 4). The subsections below outline the significance of the aforementioned sub-processes detailing the settings used in each process.

6.2.1 Step 1: Importing images to APS environment

Much like all SFM programs, APS uses a set of overlapping images of an object to reconstruct an object surface. Figure 6.1 shows the general process pipeline for the reconstruction of an object in using SFM.

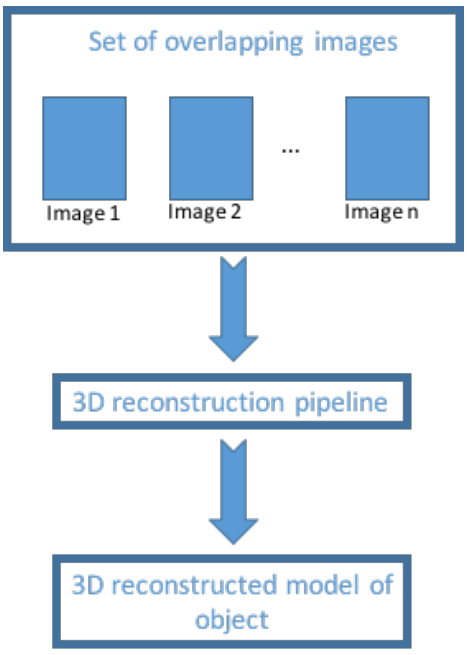

Figure 6.1: Process pipeline for photogrammetric 3D reconstruction using structure-from-motion.

The process depicted in Figure 6.1 is usually manually implemented by accessing the required images from the computer storage unit. The 3D reconstruction pipeline (Figure 6.1) consists of step 2 to step 4 of the reconstruction process.

In this project, 3D reconstruction process is automatically initiated by starting the APS application on the computer and running the Python batch processing script developed for this project within the APS environment. The images used in the reconstruction process are automatically accessed from a patient-specific folder, imported into the ASP environment, and assigned to a single "chunk" in the program. Figure 6.2 shows an illustration of 7 images inputted through by the Python script into APS and the resulting 3D reconstruction.

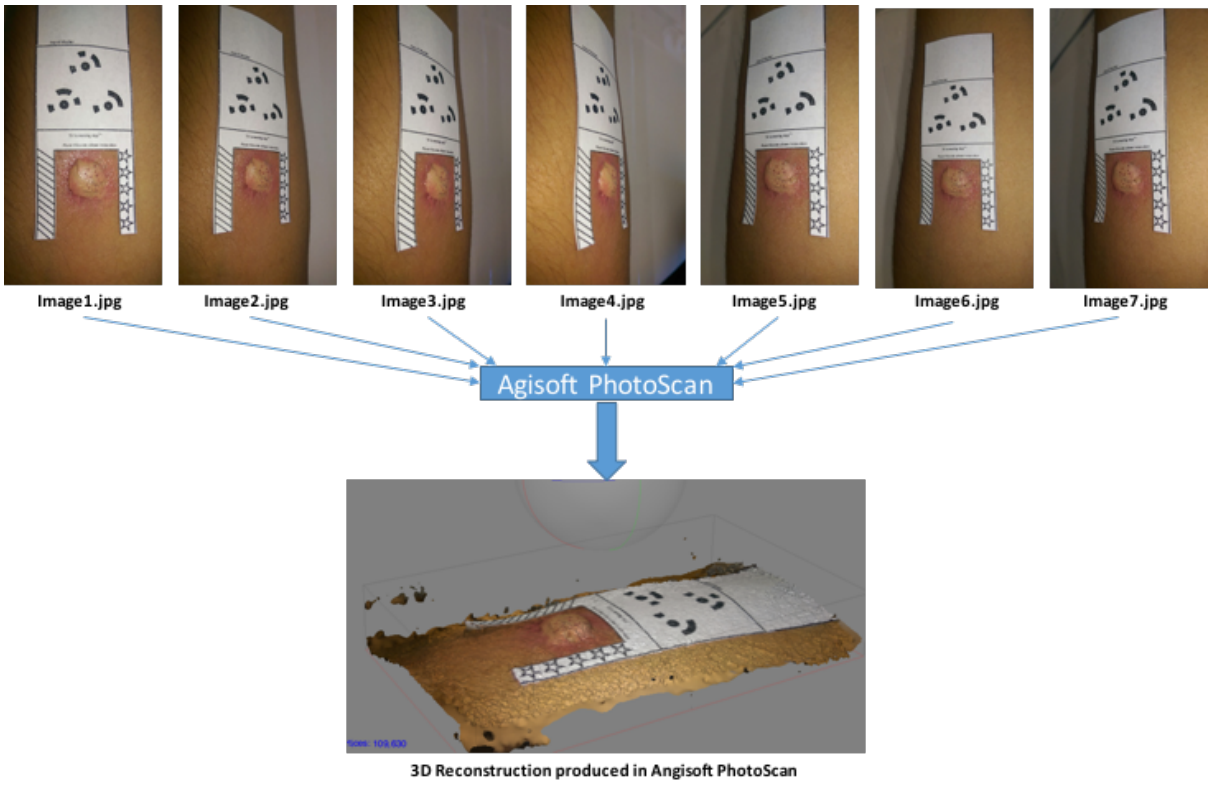

Figure 6.2: Illustration of 7 image input into APS and resulting 3D reconstruction.

The input images shown in Figure 6.2 are renamed Image1 to Image7. Image1 corresponds to the top view image, Image2 to Image4 correspond to the left side images and Image5 to Image7 correspond to the right-side images.

6.2.2 Step 2: Camera alignment and marker detection

Camera alignment is the first stage of the reconstruction process in which APS scans the input image set to find and match common points (also known as tie-points) in the images.

The set of tie-points produced from this process form a sparse point cloud; however, this point cloud is not used in further 3D model reconstruction but rather is used solely for determining image parameters, positions, and orientations in the 3D space (camera calibration parameters). A greater number of tie-points produced in the sparse point cloud results in more accurate camera calibration parameters. Furthermore, the use of coded markers and distinct surface features (which can be matched across the image set) increases the accuracy of the camera calibration parameters by increasing the number of tie-points. The images in the image set can be aligned using various settings. The 'high accuracy' setting results in a point matching process which is conducted on the image set in their original size. Lower accuracy settings downscale the images in the set thus reducing the processing time required in this phase of the reconstruction process.

For this project, although the orientation of the cameras is available in the pitch and roll data recorded by the developed application, it was found that providing this data as an input for the image orientations in APS resulted in fewer tie-points being produced. This could be due to the accuracy of the recorded pitch and roll data from the smartphone. Gyroscope data from mobile devices tend to be noisy and thus this data may not be suitable for accurate photogrammetry. As such, the specified image orientations as per the image capture protocol were only used to ensure that there was sufficient image overlap in the image set and not as an input for APS reconstruction. The additional surface features of the scaling sticker added to the flanks of the induration increased the number of feature spots which could be matched in the image set. This resulted in an increase in the number of tie-points in the produced point cloud thus increasing the accuracy of the camera calibration parameters.

Accuracy is a key component of the objectives of this project. As such, the high accuracy setting was best suited for the reconstruction process. In addition, the tie-point limit was set to 'maximum' to ensure that as many tie-points as possible were used to produce the sparse point cloud. Furthermore, to increase the accuracy of the camera calibration parameters, the coded markers were detected in this phase of processing. The distance between two coded markers was set to their real-world distance (14.5 mm). This corrects the scale of the

3D model and informs APS's estimated altitude (real world height from the arm surface) at which each image was captured.

6.2.3 Step 3: Dense cloud generation

Using the estimated image positions and orientations, APS builds a dense point cloud which estimates the surface of the object. The dense cloud is built by estimating the depth information from each camera position using the scale of the 3D model. There are two main parameters which can be adjusted in this process; quality and depth filter mode. The quality parameter specifies the desired quality (accuracy) of the reconstruction. This parameter can be set to 'high' or 'low' and is similar to the accuracy parameter of the camera alignment process in that a low-quality specification downscales the images in an effort to save processing time. APS calculates depth maps for every image in the set and uses a filtering algorithm to exclude possible depth outliers (sudden/unexpected variations in reconstructed depth) caused by bad focus and/or noisy pixels. Mild, moderate, and aggressive depth filters can be used to eliminate such outliers however, an aggressive depth filter may result in the elimination of clusters of dense cloud points that may be part of the reconstructed surface. The elimination of such clusters of points results in the formation of holes on the object surface. According to the APS user manual, the formation of holes is a common problem when dealing with small image sets (less than 10) and images of low resolution (5 MB or less).

For this project, since the image set used in the reconstruction process is both small (7 images) and of low resolution (5 megapixels), the depth filtering mode was set to mild to minimise the probability of hole formation in the dense cloud. In addition, the high-quality setting for the 'quality' parameter was used to ensure the highest accuracy possible.

6.2.4 Step 4: Mesh generation and model orientation

The final stage of the reconstruction process is the construction of a mesh that estimates the surface of the object using the dense point cloud. APS uses a 3D polygonal mesh representation to estimate the surface of the object by interpolating between points in the dense point cloud. There are three main parameters which can be adjusted in this process; surface type, polygon count and interpolation. The 'arbitrary' surface type parameter is best suited for the reconstruction of arbitrary objects while the 'height field' setting is best suited

for aerial photogrammetry (Schenk, 2005). The polygon count specifies the maximum number of polygons which will be used to represent the mesh. The optimal number of polygons in a mesh correspond to the level of surface detail required in the reconstruction process. If the number of polygons is too small, this will result in a rough surface. Conversely, a high number of polygons will produce a more detailed surface however, this will be at the expense of processing time. The final parameter, which is specified as a Boolean (either enabled or disabled), is the interpolation parameter. Disabling the interpolation parameter usually results in a mesh with holes as only the areas corresponding to high dense point cloud density are reconstructed. In this case, manual hole filling is required.

In this project, the surface type parameter was set to arbitrary as this is best suited for the reconstruction of an arm and induration because the arm surface can be likened to an arbitrary object. High detail is required for this project, thus the polygon count was set to 'high'. In addition, since the reconstruction process is automated, manual hole filling is not possible thus the interpolation was enabled for this project.

Following the mesh generation process, the 3D reconstruction was reoriented so as to complement the data extraction process. The top view image (Image1) was orientated parallel to the x-y plane and used as a reference for orientation the remaining 6 images. In this way, the z co-ordinate (in 3D space) of each point in the mesh represents the depth of that point or more precisely the relative distance of that point to the top view image. As will be seen in the next section, this information was used to produce a depth map of the arm and induration from the perspective of the top view image.

6.3 Data extraction from 3D reconstruction for automated evaluation of induration diameter

6.3.1 Scaling factor extraction and model depth map generation

Normally in APS, the reconstructed model is scaled using a reference distance on the mesh which is manually inputted. Following this, the operator can manually set marker points directly on the mesh and measure the distance between these points using the measuring

tool. Alternatively, in cases where manual manipulation of the 3D reconstruction is not possible, APS allows the user to project markers (automatically detected coded markers) on the point cloud onto any image in the image set. This proves useful for scaling images (in the 2D realm) outside the APS environment. Furthermore, APS allows the user to produce additional 2D images (in addition to the original image set) which can be constructed to represent various properties of the dense cloud points. One such 2D image is a depth map of the 3D reconstruction from a specified perspective (position and orientation in the 3D space). Depth maps like these, which represent the change in depth of the surface using colour variations over the produced image, could be used in a post-processing pipeline to identify height-specific features of the 3D reconstructed surface.

The common method of manually measuring features on the 3D reconstruction is not suited for an automated distance measurement process in APS as markers cannot be assigned to point cloud data points using Python scripting. As such, in this application, the process of identifying, scaling, and measuring the induration's critical diameter is conducted outside the APS environment using a depth map produced from the 3D reconstruction.

The 3 markers on the scaling sticker are detected in the APS 3D environment and the locations of these markers are projected back on the top view image. The marker locations, which are recorded in a text file, are represented by the pixel location of each marker in the top view image. The Euclidian distance between two of the markers is calculated using Equation 1 below:

$$Marker\ distance = \sqrt{(x_{marker1} - x_{marker2})^2 + (y_{marker1} - y_{marker2})^2} \qquad \text{[Eq. 1]}$$

where the marker distance is in pixels and the ***x*** and ***y*** values of the marker coordinates represent the pixel locations of two of the 3 markers on the top view image. Since the markers are arranged in the form an equilateral triangle, the selection of the two markers used in Equation 1 was arbitrary. Using the pixel distance and the real-world distance between the two markers, a scaling factor (in millimetres per pixel) is determined by

dividing the real-world distance (in millimetres) by the marker distance (in pixels). This value is recorded in a text file and stored in a patient-specific folder.

6.3.2 Depth map generation

The top view image is a central part of the evaluation process and thus most of the data obtained from the 3D reconstruction is extracted with reference to this image. Since the model is orientated such that the top view image is parallel to the ***x-y*** plane in the APS 3D environment, the depth (height difference) of every point in the mesh is represented by the z-coordinate value of each point in the dense cloud. Figure 6.3 shows a visual representation of this in the APS environment.

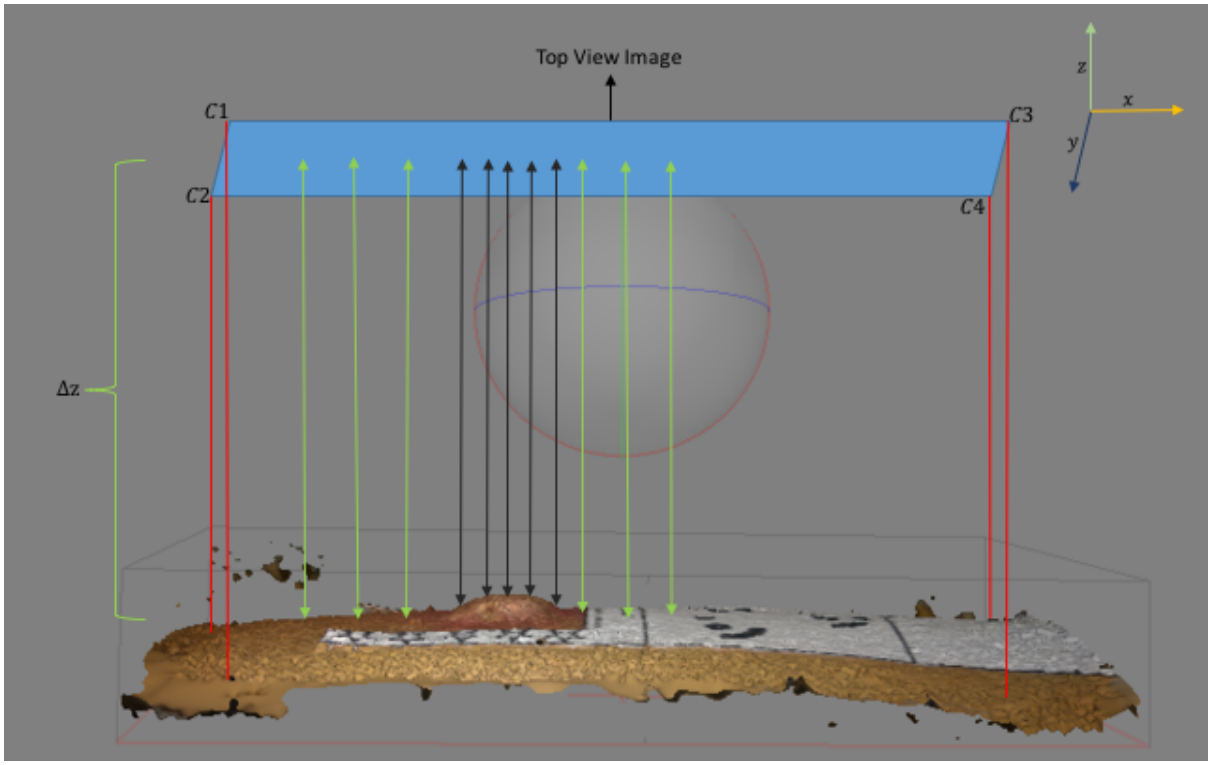

Figure 6.3: Depth map generation using z-coordinate of surface mesh in APS 3D environment.

Figure 6.3 shows the top view image above the 3D reconstruction and parallel to the ***x-y*** plane with corners C1, C2, C3 and C4. The corners pixels (C1 to C4) of the image are projected onto the 3D mesh as illustrated by the red lines in Figure 6.3. The rectangular area on the mesh formed by these four projections is the enclosed area of the mesh which encompasses all features which were captured in the top view image. Following this, every pixel in the image is projected onto the mesh and the ***z***-coordinate of the point on the mesh corresponding to each pixel of the top view image is recorded. Since the top view image is required to be perpendicular to the surface of the forearm (as per the stipulated pitch and

roll specification in the image capture protocol), the z-coordinate of the points represent the height of the surface of the skin at that point – a lower Δz value means the skin higher. The Δz values of a series of points running over the induration are shown in Figure 6.3 with the black and green double arrows representing points on the induration and points on the surrounding skin, respectively. The minimum and maximum values of the set z-coordinates are recorded and mapped to a scale of 0 to 255 (corresponding to grey-scale intensity range), with the maximum value mapped to 0 and the minimum value mapped to 255. All other points in the set of z-values are mapped to values between 0 and 255 accordingly.

Using the transformed map of points between 0 and 255 as grey-scale pixel intensities, a grey-scale image – a depth map -, with dimensions matching the top view image, is created. Figure 6.4 shows an example of such a grey-scale depth map image generated using the scaled depth data alongside the original top view image of the induration.

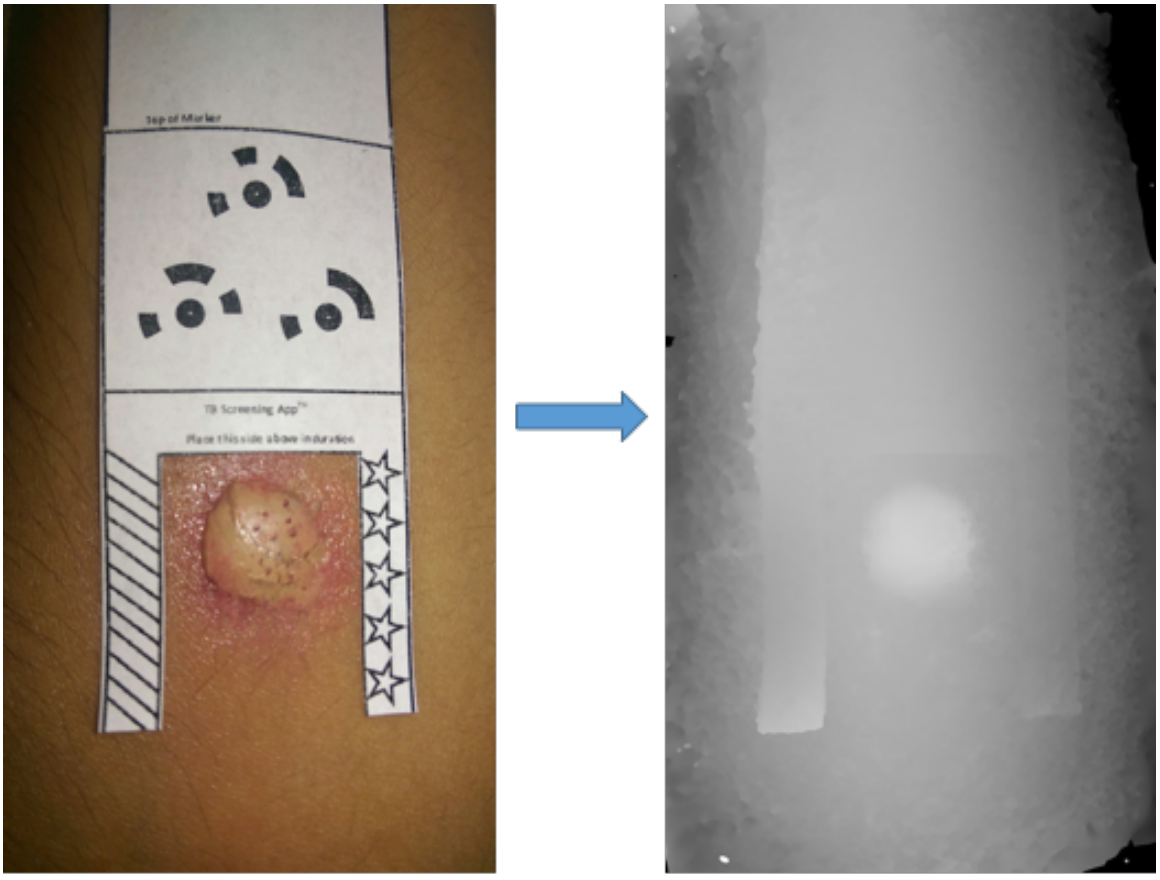

Top view image

Depth map

Figure 6.4: Original image and depth map of 3D reconstructed surface from top view image perspective.

As seen in Figure 6.4, raised areas (the induration and parts of the scaling sticker) are lighter than the lower areas (the surface of the skin). The pixel dimensions of the generated depth map image are equal to that of top view image (in this example, 1520 X 2688 pixels) so that the position of features in the original image correspond to the position of the same features in the depth map image. Since the location of the induration and orientation of the long axis of the arm are specified in the image capture protocol, these features are directly identifiable on the depth map image.

6.4 Database design and implementation

Database construction and management play a pivotal role in organising and storing results of the reconstruction and evaluation processes. The inputs (original images), intermediate results (depth map and 3D model) and final results (diameter of the induration) of the process must be organised and stored both for access by the automated evaluation algorithm and for review by a trained professional.

Since Python scripting is supported as an interface to APS for batch processing, the automated construction of the database was also coded in Python. This allowed for a seamless integration between database construction for each patient and initiation of the automated reconstruction process in APS.

During real-world operation, the output package from the mobile application is downloaded to the remote computer and used as the input for initiating the evaluation process (3D reconstruction and measurement of induration diameter). The files contained in the output package are extracted into various folders in a “Main Data Folder” on the computer. The Main Data Folder contains folders for each patient (patient-specific subfolder) where intermediate and final results of the evaluation process are stored for later access. Figure 6.5 shows the subfolders contained in each patient folder and how the package contents and intermediate results of the reconstruction process are distributed in these patient-specific subfolders.

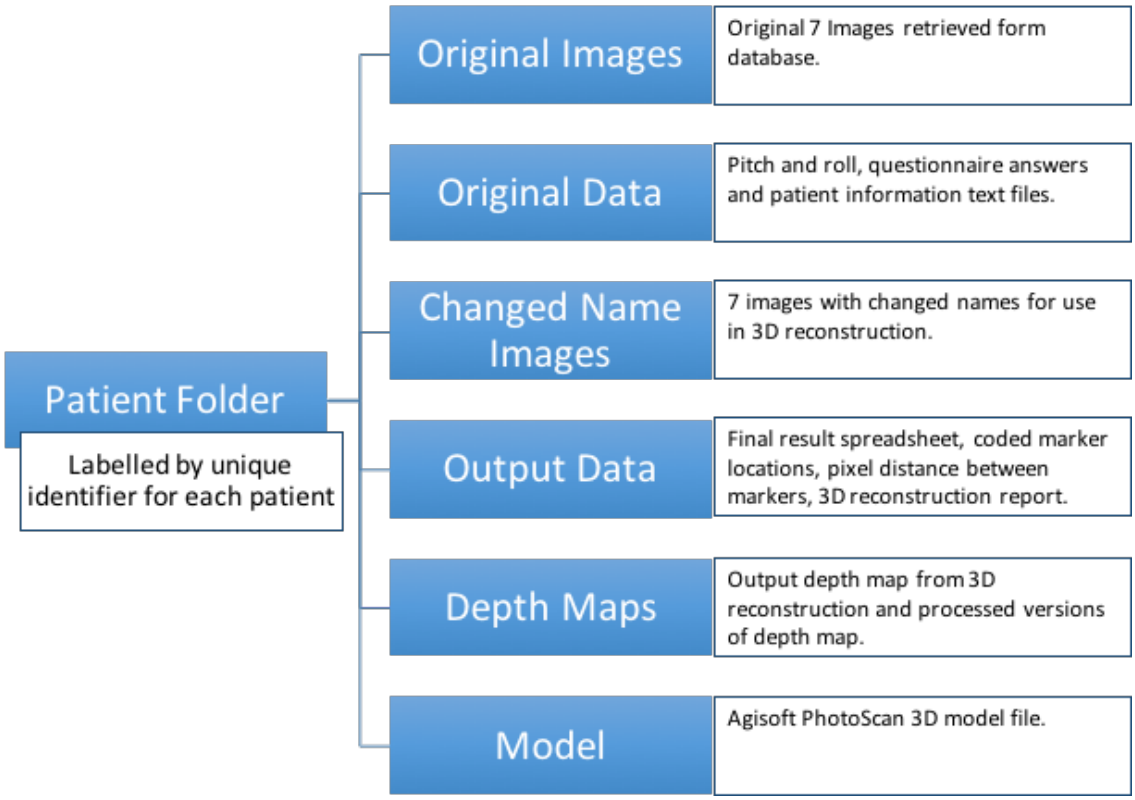

Figure 6.5: Patient subfolders and subfolder contents in the database.

Each patient folder contains six subfolders (Figure 6.5). Upon initiation of the Python script, the algorithm creates a patient folder with the subfolders for each patient detected in the online database. Each patient is assigned a unique identifier to keep track of patient data. A description of the contents and rationale for each subfolder is as follows:

- Original Images: Contains the original images with names as labelled by the mobile device upon capturing the images. These images are simply copied into this folder and used as a source of the original image data.
- Original Data: Contains the three text files: pitch and roll data, questionnaire answers, and patient information. The contents of this folder are inputs for the automated evaluation algorithm.
- Changed Name Images: The 7 images contained in this folder are copied from the "Original Images" folder, however, the names of the images are changed so that they may be easily accessed and referenced by the reconstruction and evaluation algorithms.
- Output Data: Contains various intermediate results from the reconstruction and evaluation processes. The spreadsheet contained in this folder records a

summary of the critical aspects of the processing pipeline and result of the evaluation (i.e. size of the measured induration).

- Depth Maps: Contains the depth map produced from the reconstruction as well as the processed depth maps used in identifying the induration and measuring the result.
- Model: Contains the APS 3D reconstruction of the induration and the arm. This allows clinicians to re-evaluate the induration in the 3D space.

The algorithm constructs the abovementioned folders for a single patient and imports the relevant data to each folder before initiating the reconstruction process for the patient.

7. AUTOMATED MEASUREMENT OF INDURATION FROM 3D RECONSTRUCTION

The size of the induration was measured outside the APS environment using the depth map and marker information extracted from the 3D reconstruction process. This chapter presents an overview of the automated measurement of the induration, followed by a detailed description of each phase. The process of measuring the induration has various sub-processes which result in the final output (size of induration). Figure 7.1 shows a process flow diagram.

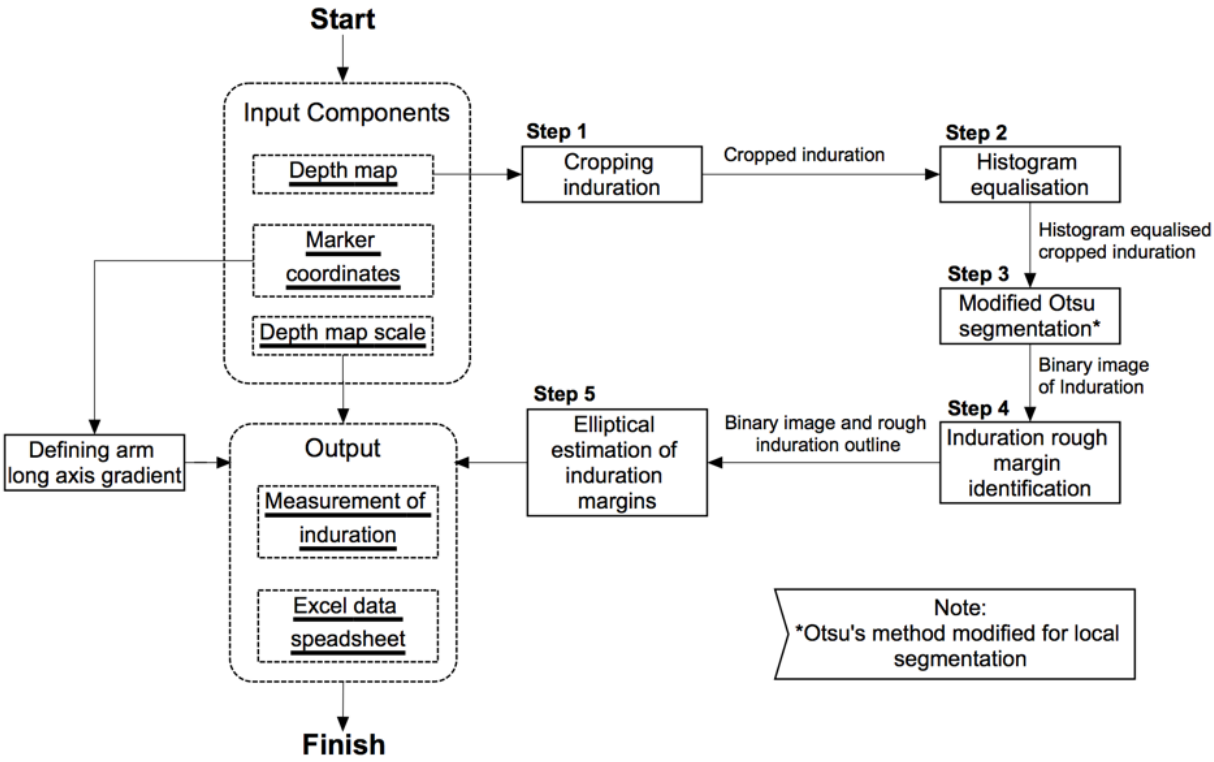

Figure 7.1: Flow chart of automated induration measurement using results from the 3D reconstruction.

As shown in Figure 7.1, three input components were used in the automated measurement process: the raw depth map, pixel coordinates of the 3 markers and the scale of the depth map in pixels per millimetre. The depth map image was passed through an image processing chain to identify the induration margins. The final output of this processing chain, combined with the orientation of the long axis of the arm and scale of the depth map, were used to measure the size of the induration. The details of each of the steps in the image processing chain shown in Figure 7.1 are outlined in subsection 7.2.

7.1 Feature matching

The depth map generated from the view of the top view image has the same pixel dimensions as the depth map image. Thus, the features and areas in the top view image can be directly mapped to locations on the depth map image. Figure 7.2 shows a set of features from the top view image (on the app image capture screen) that are mapped to the depth map image.

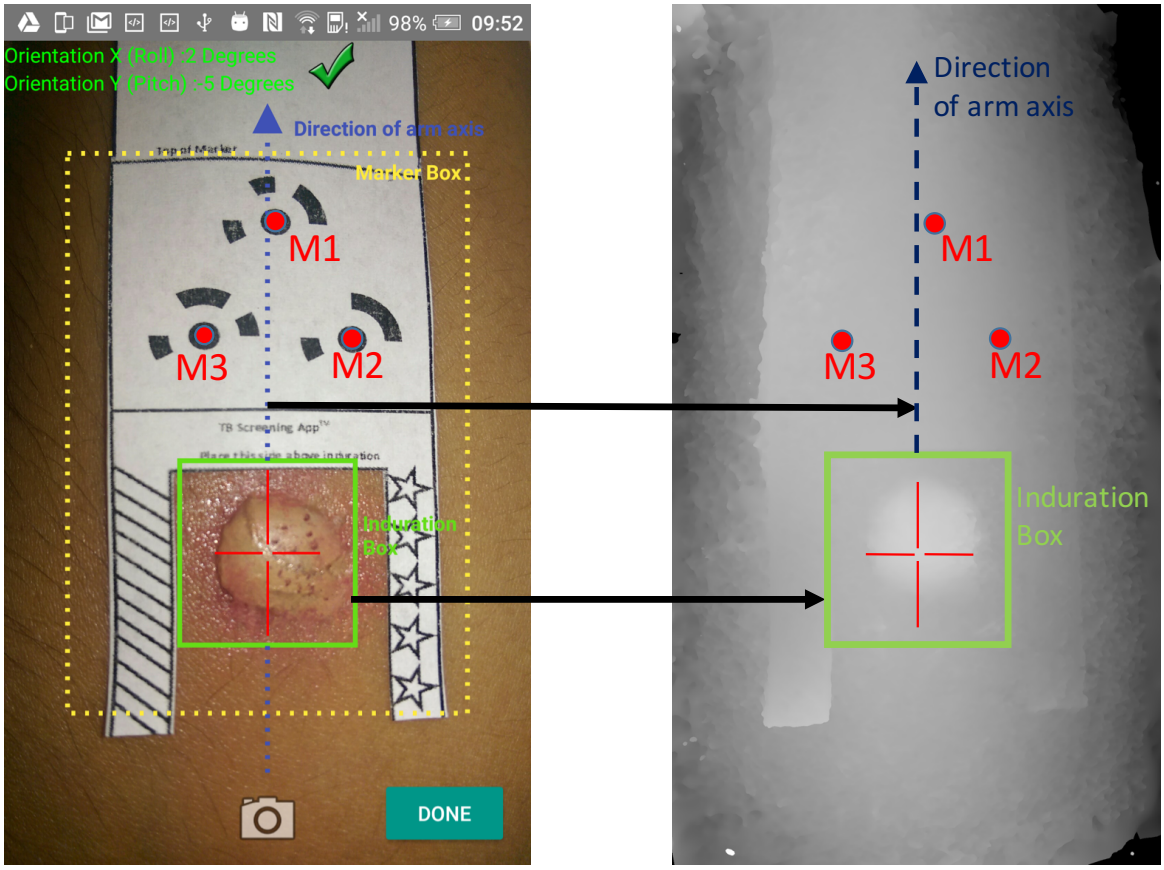

Figure 7.2: Feature matching between the top view image and the depth map image.

As illustrated in Figure 7.2, three components needed for measurement are mapped from the top view image to the depth map image: the induration box, direction of the capture screen long axis of the arm and the positions of markers M1, M2 and M3. The induration width is measured perpendicular to the long axis. The axis shown by the blue dashed line runs vertically upward on the image capture screen and is mapped to run vertically upward (parallel to the long sides of the of the image) on the depth map image. In the subsections to come, a second method of finding the scaling sticker arm long axis line using markers M1, M2 and M3 will be presented.

7.2 Induration identification

7.2.1 Step 1: Cropping of depth map image

The first stage of the evaluation process is cropping the induration box area of the depth map. This is done by mapping the pixel coordinates of the induration box on the capture screen (of the top view image) to the pixel coordinates of the same box on the depth map image and finally cropping this area. Figure 7.3 below illustrates the area matching process from top view image to the depth map image and the resulting cropped induration image.

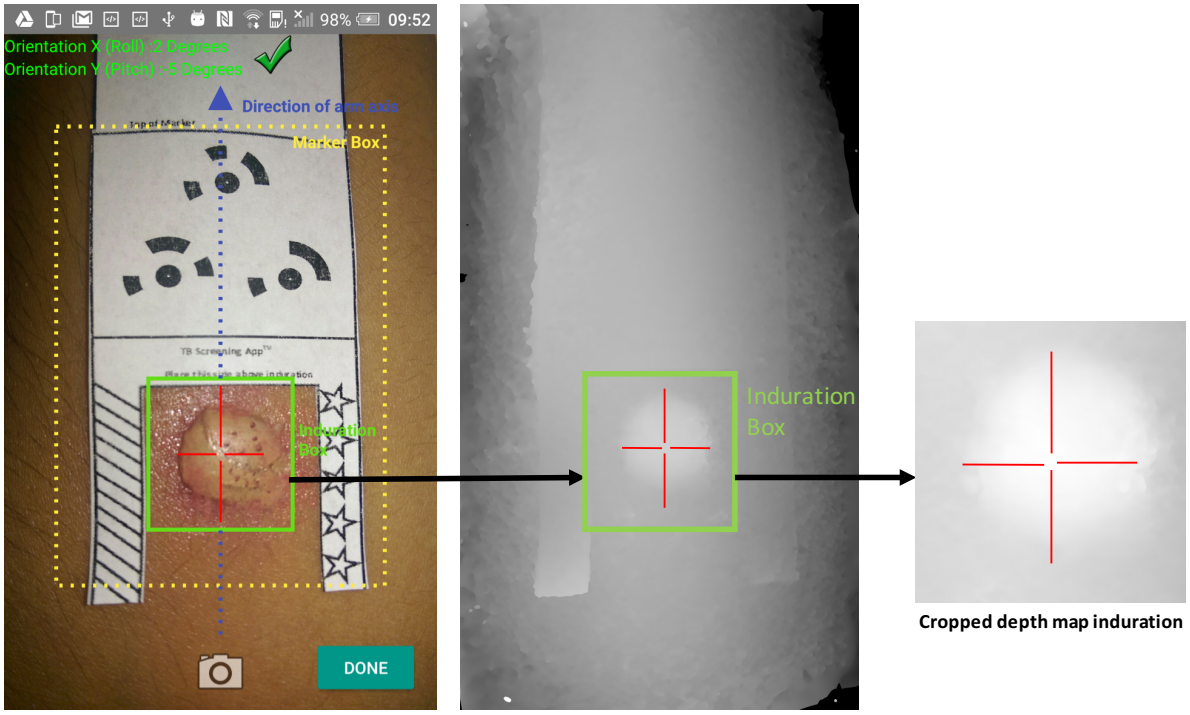

Figure 7.3: Induration cropping step on the depth map using induration box area mapped from the top view image.

The image capture protocol and image capture specifications for the top view image (see Section 4.1.4) are designed to increase the likelihood that the entire induration is positioned in the induration box with the induration covering as much of the induration box as possible. A second constraint on the operator (when capturing the top view image) is that the red induration box target (see Figure 7.3) be placed on the induration. These two constraints ensure that most of the cropped image (from the depth map image) contains the induration and that the centre of the cropped image (the red target) always represents a pixel which can be identified as part of the induration.

7.2.2 Step 2: Histogram equalisation of cropped induration image

Histogram equalisation is a widely used method to enhance contrast in images (Kim, 1997). Objects or features in low contrast images may be difficult to identify and applying a histogram equalisation algorithm can assist in clarifying boundaries and features of objects in an image. By flattening the density distribution of pixels in an image the overall contrast of the image is increased; this is commonly known as stretching the dynamic range of the image (Kim, 1997). Histogram equalisation has been shown to improve the performance of unsupervised bimodal segmentation and thus object identification in greyscale images (Sklansky, 1978).

The second phase of the processing pipeline is concerned with applying a histogram equalisation algorithm to the cropped induration image. The cropped image of the induration shown in Figure 7.3 has low contrast between the induration and the surrounding skin. This is expected as the depth data (over the entire depth map) extracted from the 3D reconstruction is spread over the 0 to 255 grey scale. As the induration is a small height feature in relation to the depth change of the surface of the arm (curved arm surface), the depth change between the induration and the skin immediately surrounding it is allocated a small range of the grey scale. Figure 7.4 shows the result of applying histogram equalisation to the cropped induration image.

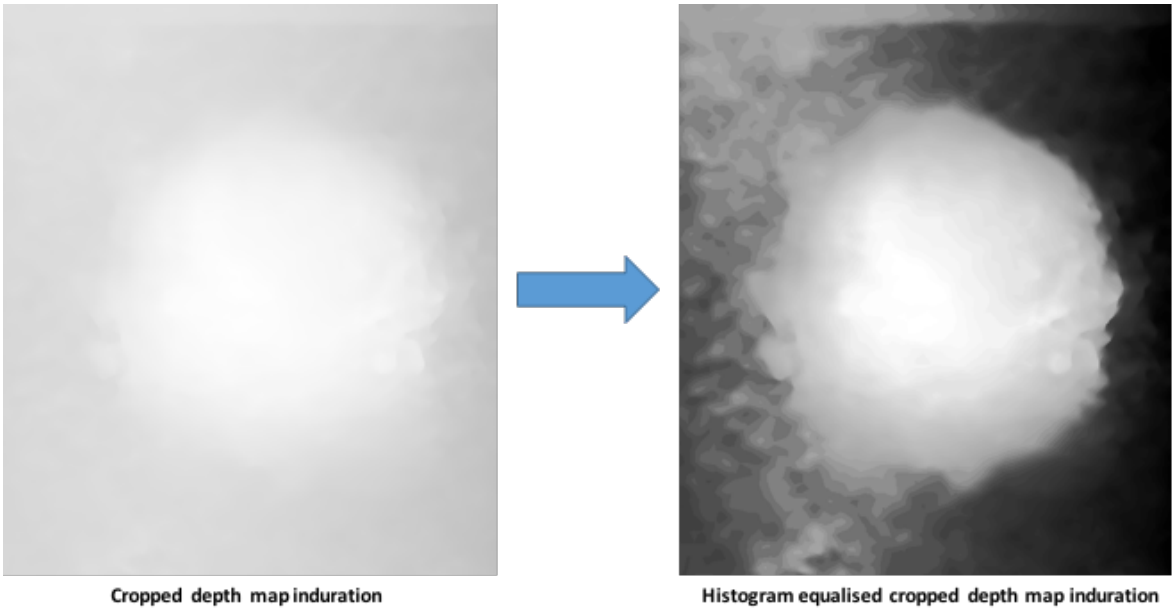

Figure 7.4: Histogram equalised applied to cropped depth map induration image.

As shown in Figure 7.4, the contrast of the induration relative to the surrounding skin has been increased. This has a twofold purpose: firstly, it can be used by a clinician at the remote processing unit to visualise the induration more clearly for reassessment of the result. Secondly, it improves the performance of the segmentation method used to identify the margins of the induration. The latter purpose will be examined further in the next subsection.

7.2.3 Step 3: Segmentation of cropped induration image

7.2.3.1 Bimodal segmentation: Otsu's method

Bimodal segmentation is used extensively in image processing and some computer vision applications for identifying object boundaries in greyscale images containing two objects represented by two pixel classes, respectively (Otsu, 1975). Otsu (1975) suggested a nonparametric and unsupervised method of selecting the threshold value for bimodal segmentation. Since the image to be segmented has two distinct pixel classes (foreground and background), Otsu's method assumes that the pixels in the image follow a bi-modal histogram distribution. Figure 7.5 shows a general bimodal representation of pixel intensities in an image containing two pixel classes.

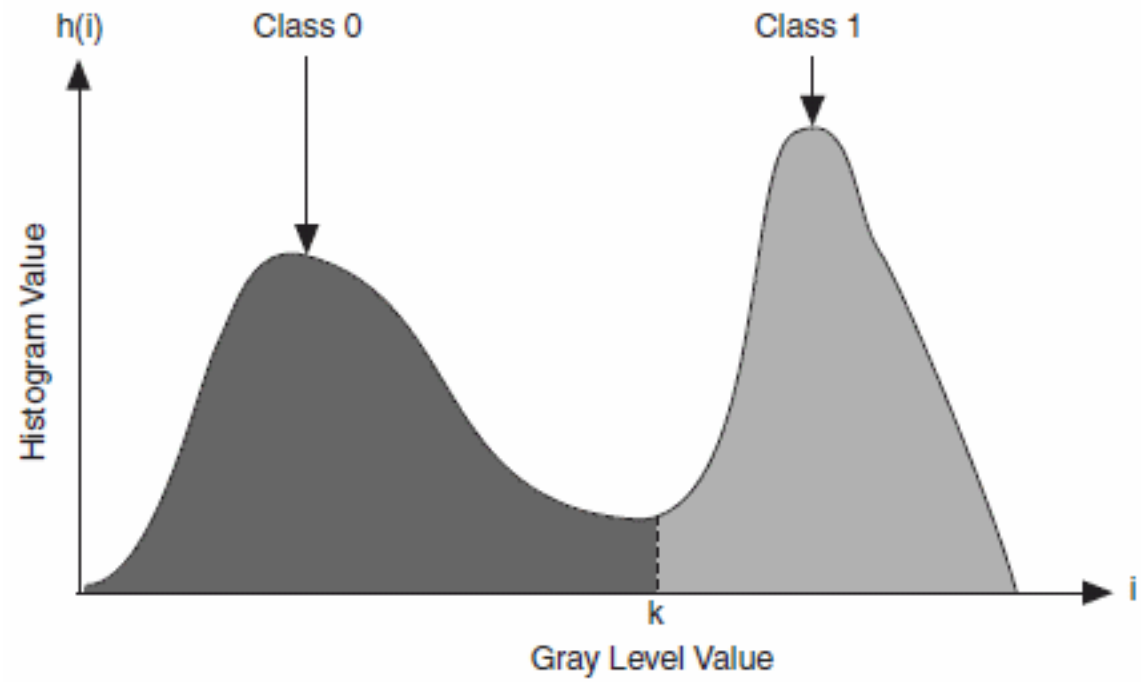

Figure 7.5: Typical bimodal greyscale pixel intensity histogram showing two distinct classes (image reproduced from Instruments (2012)).

In Figure 7.5, h(i) represents the frequency of occurrence of grey level values on the x-axis of the graph. The segmentation threshold k in Figure 7.5 is determined by an iterative method which selects a value of k that results in minimal intra-class variance in each of the two classes. The primary advantages of this method are that the algorithm requires little computing power and can set a segmentation threshold value based on information in the image.

The cropped image of the induration can be treated as a bimodal grey scale image. This means that the image can be assumed to contain two "objects" (induration and surrounding skin) which can be classified as foreground and background respectively. Thus, a bimodal segmentation method can be used to identify the induration and segment it from the surrounding skin. However, since this process is fully automated and different depth maps could have unique thresholds for segmentation, a segmentation threshold value cannot be manually specified. As such, segmentation by Otsu's method was well suited to the problem. Figure 7.6 shows an example of Otsu's method used to segment a cropped induration depth map of a simulated 3D printed induration.

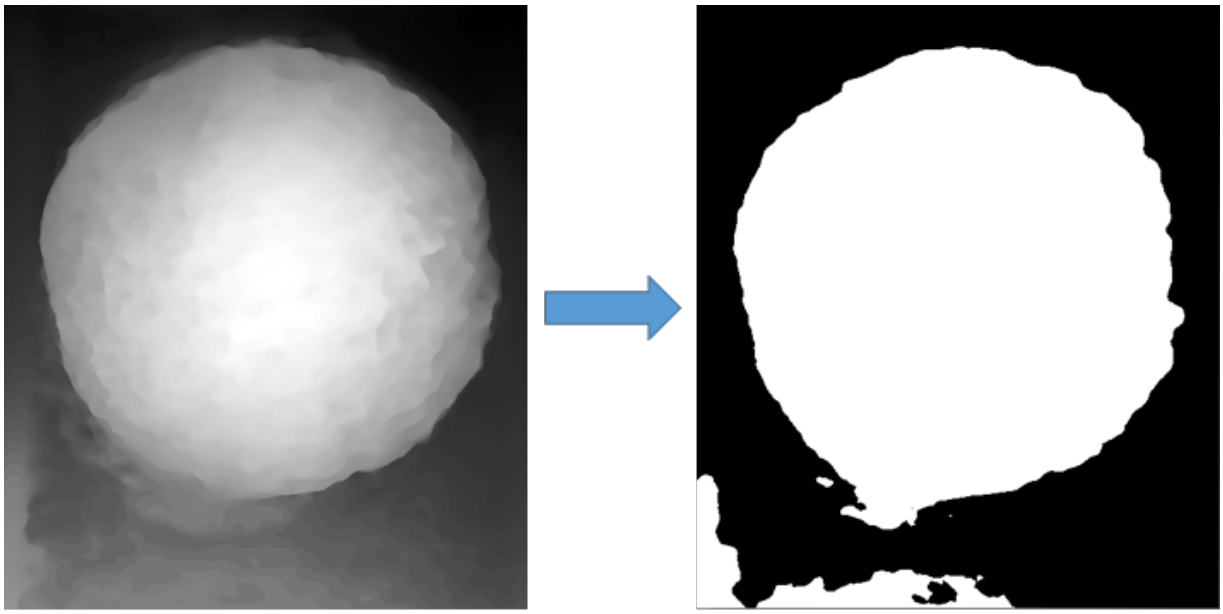

Figure 7.6: Example of Otsu's method segmentation on depth map of a 3D printed induration.

Otsu's method works well in the case depicted in Figure 7.6 as it identifies the margins of the induration clearly. However, as seen in the bottom left corner of the segmented image, the algorithm has identified an area that is clearly skin (by visual inspection) as part of the induration. This error can be termed as an "unconnected misclassification error" as the incorrectly segmented pixels are not connected to the correctly segmented object (the induration). This result occurs because the pixels in the bottom left area of the histogram equalised cropped image are lighter than the typical pixel which would be identified as skin as defined by the threshold value determined using Otsu's method. This could be because of noise in the depth reconstruction (during 3D reconstruction the induration) which results in areas of the skin which appear to be more elevated than they are. Such noise in the depth reconstruction is a common occurrence caused by a lack sufficient depth information from the input image set. However, this effect is expected as the input image set was small and of low resolution. In the case presented in Figure 7.6, this error in segmentation can be rectified with further image processing because the area identified as the induration is not connected to the incorrectly segmented area in the bottom left corner.

However, noise in the depth data is not the only cause of segmentation error. The example histogram equalised depth map illustrated in Figure 7.6 shows a well-defined induration (clear and consistent grey-scale contrast between the induration and surrounding skin on the induration edges). However, in some instances, the TST indurations could be shallow and unsymmetrically elevated from the skin. In these instances, there may not be a clear general contrast between the induration area and surrounding skin across the entire image. Thus, a single segmentation threshold value cannot be used to define the area which represents the induration. In such cases, the pixels in the depth map image representing the surrounding skin may have similar grey scale values to those representing the induration in localised areas around the induration. As such, the segmentation method could fail by incorrectly extending the segmented induration to areas which should represent the surrounding skin (connected misclassification error). Figure 7.7 shows an example where Otsu's method fails to segment the induration correctly.

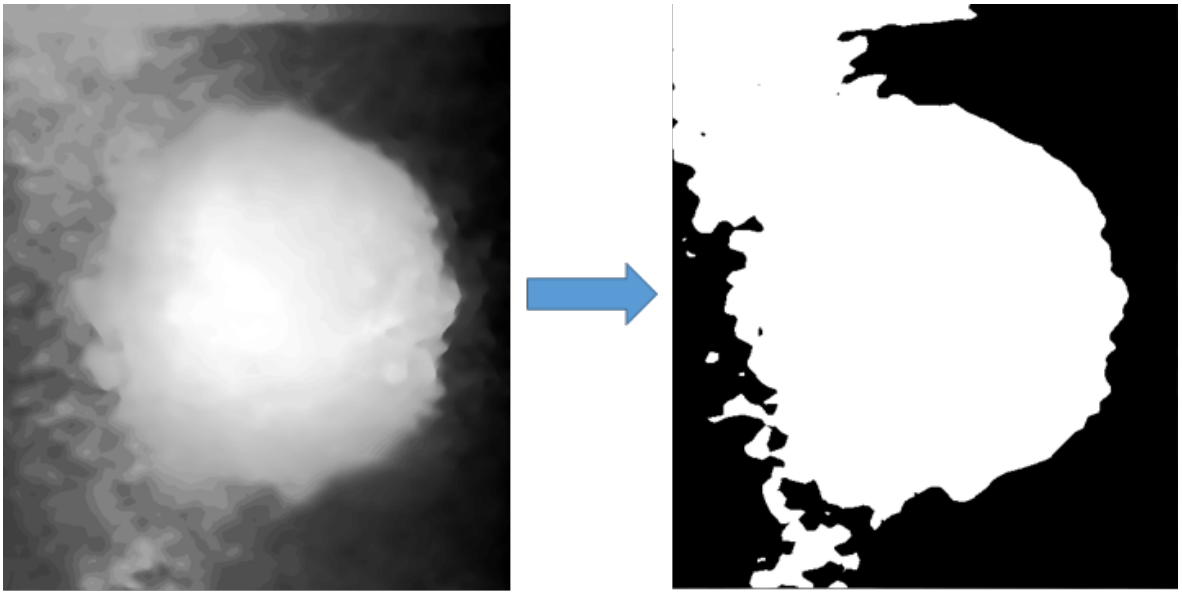

Figure 7.7: Example of failed segmentation (connected misclassification error) using Otsu's method.

The segmented induration area is incorrectly extended to the top left corner of the image. This error cannot be rectified with further processing as the incorrectly segmented area is connected to the correctly segmented object and thus the incorrect segmentation cannot be eliminated without objectively identifying a cut-off point for the induration margin.

Otsu's method uses pixel intensity information from the entire image to decide on a threshold segmentation value, which is used to distinguish between induration and surrounding skin. This causes an error in segmentation in two cases: where depth data noise is present, and where the imaged induration is shallow and unsymmetrically elevated from the skin. Both these cases could result in an undesired connected misclassification error. It is visible in the histogram equalised depth map in Figure 7.7 that the distinctions between skin and induration pixels are clear in localised areas, however, one cannot accurately assign a threshold value that would represent a cut off point for all skin pixels in the image.

7.2.3.2 A variation on Otsu's method

Moghaddam and Cheriet (2012), presented a variation of Otsu's method which varies the selected threshold based on information in a focus window (local thresholding) to better estimate background pixels to clarify blurry text on a white page. It was found in this project that Otsu's method fails to correctly segment the induration in cases where the grey-scale

intensity of skin pixels in the image vary vastly across the image. Thus, implementing a localised variation of Otsu's method may prove to be useful in this case. Figure 7.8 shows a visual representation of the approach used in modifying Otsu's method to incorporate local thresholding considering inherent properties (defined by the image acquisition protocol) of the cropped induration image.

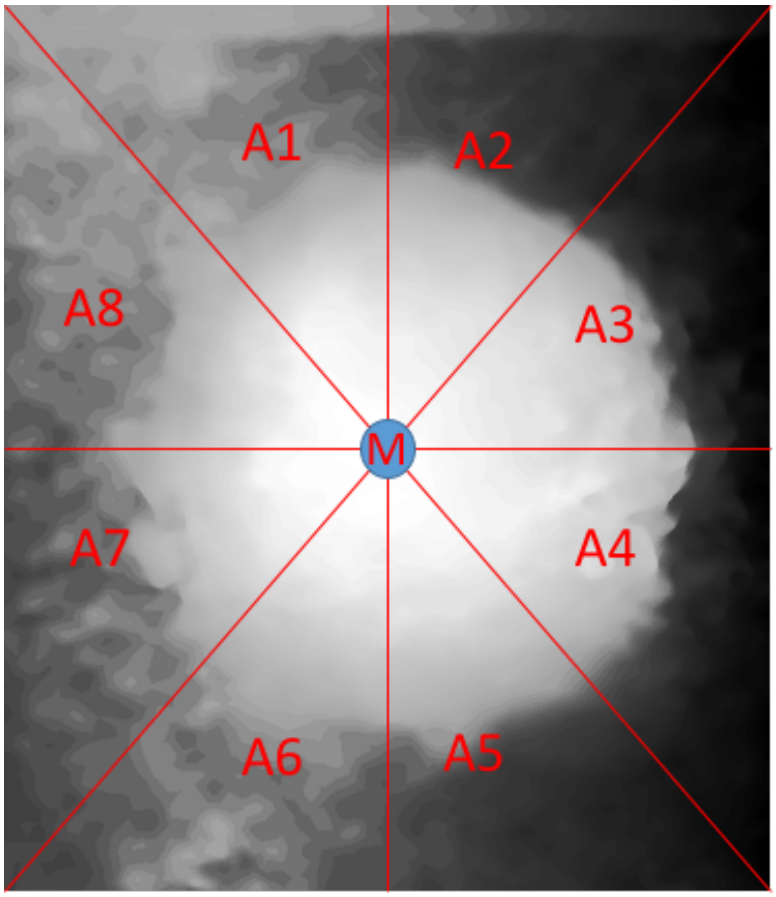

Figure 7.8: Image segments for modified version Otsu's segmentation method, to identify local margins.

Figure 7.8 shows the histogram equalised depth map image of the induration which is subdivided into 8 triangular sections represented by areas A1 to A8. Each of these sections has one of their vertices on the midpoint of the image (M) and the two other vertices on the boundaries of the image. Since the midpoint M coincides with the centre of the induration box target (red target on image capture screen) shown in Figure 7.3, one of the vertices of each triangular area lies on a point which is definitely part of the induration. In addition, since the induration does not touch the borders of the cropped induration image (the induration box) as specified in the image acquisition protocol, it can be concluded that an edge of the induration exists in each of the subdivided areas. Using this information, Otsu's

method was implemented to define the edge of the induration local to each area (A1 to A8). This was conducted by mathematically defining each area on the pixel Cartesian plane and clustering the pixels intensities into their appropriate area subdivision. Otsu's segmentation method was applied to each of these subdivided pixel clusters (A1 to A8) and a threshold segmentation value was established for each subdivision which was used to determine the margin of the induration locally. Figure 7.9 shows the result of using the modified segmentation method on the same depth map image shown in Figure 7.7.

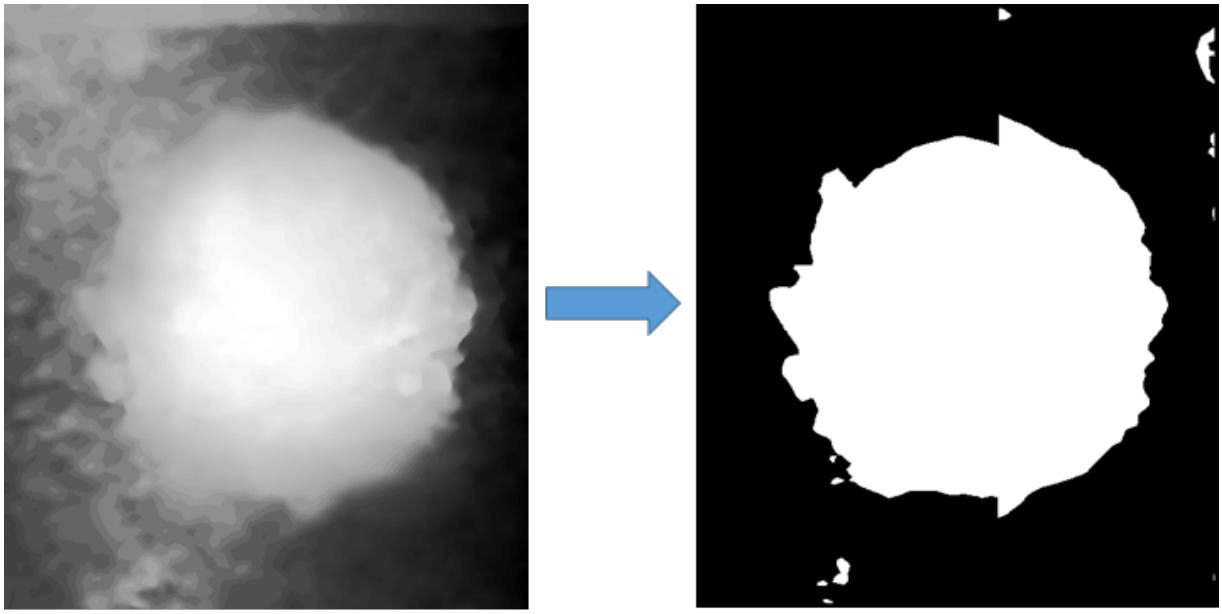

Figure 7.9: Example of segmentation using modified Otsu's method.

As seen in Figure 7.9, using the modified Otsu's method to segment the image locally results in an improved segmentation in cases where a clear and consistent contrast between induration and skin across the entire image does not exist. This algorithm was tested on multiple depth maps similar to the one in Figure 7.9 and the results were consistently better than the unmodified Otsu's method segmentation. The segmented image (Figure 7.9) still contains small patches of incorrectly classifies pixels, however, catastrophic failure (unconnected misclassification error) of the segmentation method was avoided and is less likely to occur using the modified Otsu's method.

7.2.3.3 Impact of histogram equalisation on segmentation

Although histogram equalisation produced a more visually appealing representation of the induration by increasing the contrast of the image, it was necessary to assess its impact on

the segmentation process. The impact of histogram equalisation on segmentation was assessed using the following methodology:

- Using the top view original image of a simulated skin induration, the margins of the induration were visually identified and manually segmented from the background. This was used as a ground truth segmentation of the induration.
- Segmentation using the modified Otsu's method was performed for both the histogram equalised and non-histogram equalised depth map images from the same simulated induration above.
- The equalised and non-equalised segmented images were compared to the ground truth using two quantitative measures; image structural similarity (SSIM) index and mean square error (MSE).

SSIM is used as a measure of the similarity between two images. The method is usually implemented using an image which has been distorted by some image process and the original undistorted image. The quantitative result of the SSIM gives a measure of the effect of the distortion of the image. The SSIM index is calculated by iterating through windows of size $N \times N$ on the original and distorted images and comparing the two windows by three measures: structure, luminance and contrast. However, since luminance is not a factor in this assessment and contrast difference can be independently assessed (using the MSE), only the structure comparison is of interest in this case. The SSIM index is represented on a scale of 0 to 1, with an SSIM value of 1 representing perfect structural similarity between the images being assessed.

In the context of this test, the MSE is a measure of the deviation in grey scale value between a pixel in the ground truth image and a pixel in the same position on the modified Otsu's method segmented image. The MSE is calculated by iterating through the pixels of each of the two images and summing the error between the pixels intensities. This error is squared and divided by the total number of pixels in the image. In the case of binarised images, the MSE gives a measure of how many pixels have been incorrectly segmented by the segmentation algorithm. A higher MSE value represents a higher number of incorrectly segmented pixels.

The SSIM and MSE tests were conducted on two test cases which represent the performance extremes of the segmentation algorithm: a good segmentation result and a bad segmentation result. This selection was made by visually inspecting and comparing the segmented images to images of the induration. The bad segmentation result chosen for this test was one of the simulated skin indurations with a low induration height and bad induration edge definition. By conducting this test on the two extreme test cases, it is assumed that all other improvements because of histogram equalisation can be mapped between these two extremes. Figure 7.10 and Figure 7.11 show the outcomes of the tests for a good and bad segmentation results respectively.

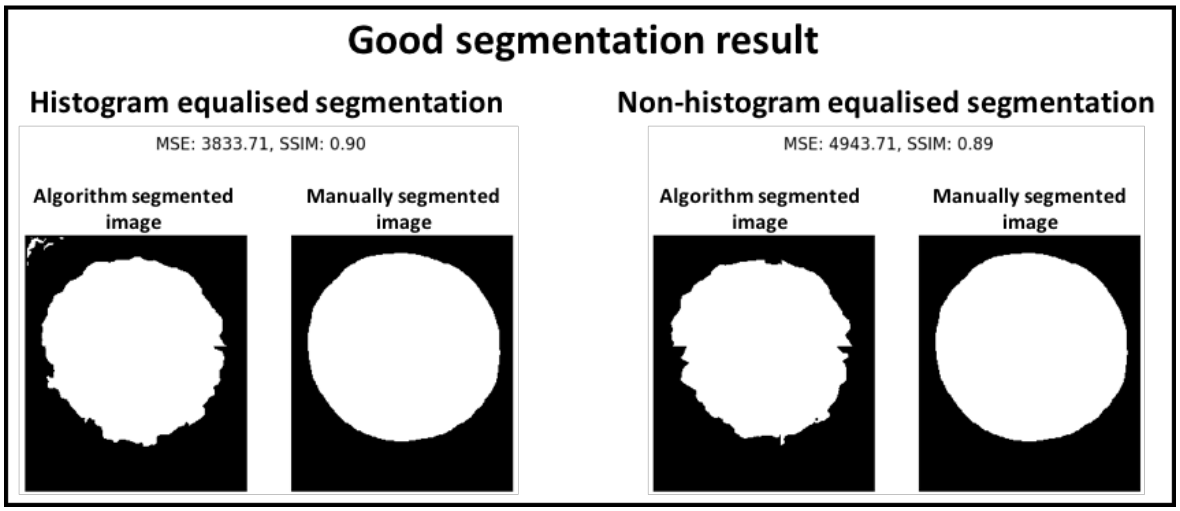

Figure 7.10: Good segmentation MSE and SSIM comparison for histogram equalised and non-histogram equalised depth maps.

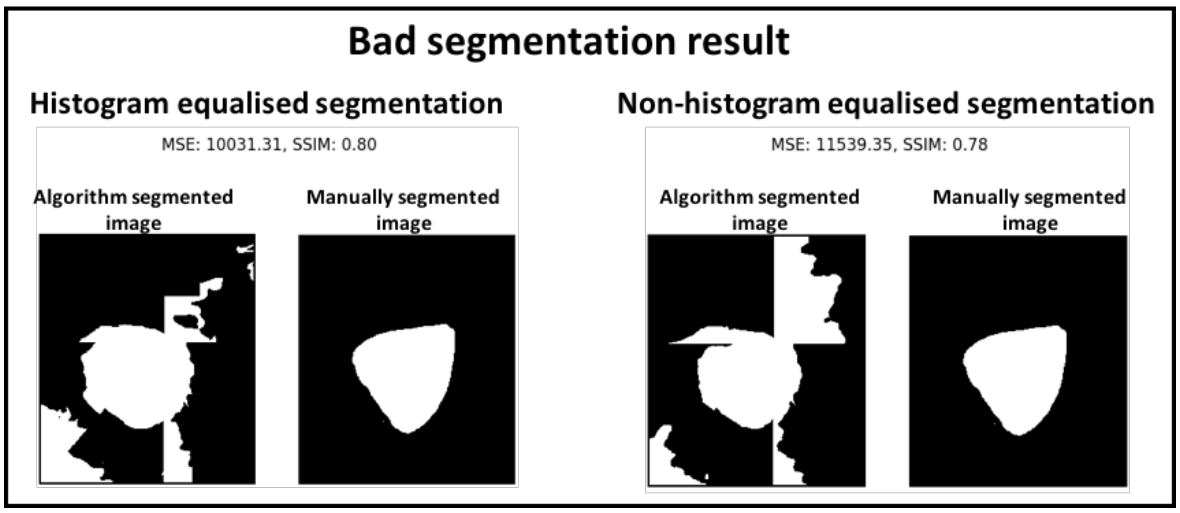

Figure 7.11: Bad segmentation MSE and SSIM comparison for histogram equalised and non-histogram equalised depth maps.

The histogram equalised segmentation had a lower MSE and higher SSIM for both the bad and good segmentation results (Figure 7.10 and Figure 7.11). Although this is not a comprehensive test of the impact of histogram equalisation on segmentation using the modified Otsu's method, the results indicated that histogram equalisation may be beneficial to the image processing pipeline.

7.2.4 Step 4: Identification of rough margin of induration

As previously mentioned, and seen in the segmented images in Figure 7.6 and Figure 7.7, the segmentation process could misclassify pixels as part of the induration when in actual fact they should be classified as surrounding skin. This problem presents itself in two ways: as unconnected and connected misclassification errors. In the occurrence of these errors the margin or area coverage (of the image) which represents the induration is distorted. Figure 7.12 shows examples of both the segmentation errors.

Figure 7.12: Examples of connected and unconnected misclassification errors from the segmentation process.

Connected misclassification errors cannot be rectified post-segmentation as this process requires an objective identification of the margin of the induration, which is not possible in an automated processing pipeline. This error was mitigated as far as possible by the modified version of Otsu's method (using local thresholding). The problem of unconnected misclassification, however, can be addressed post-segmentation.

Following the segmentation of the depth map, a process of identifying the rough margins of the induration and eliminating all unconnected misclassified clusters is implemented. This is executed by:

- Using an edge detection algorithm (OpenCV2 Python package: www.openCV.org) to outline all "object" edges in the image. This detects the edges of all unconnected misclassifications and the edge of the induration.
- Running a contour detection algorithm (OpenCV2 Python package: www.openCV.org) over these edges. Only closed contours are selected and sorted by length of the contours in descending order in an array.
- The largest contour is selected as the contour surrounding the induration.

Figure 7.13 shows an example of this process implemented on an image with unconnected misclassifications.

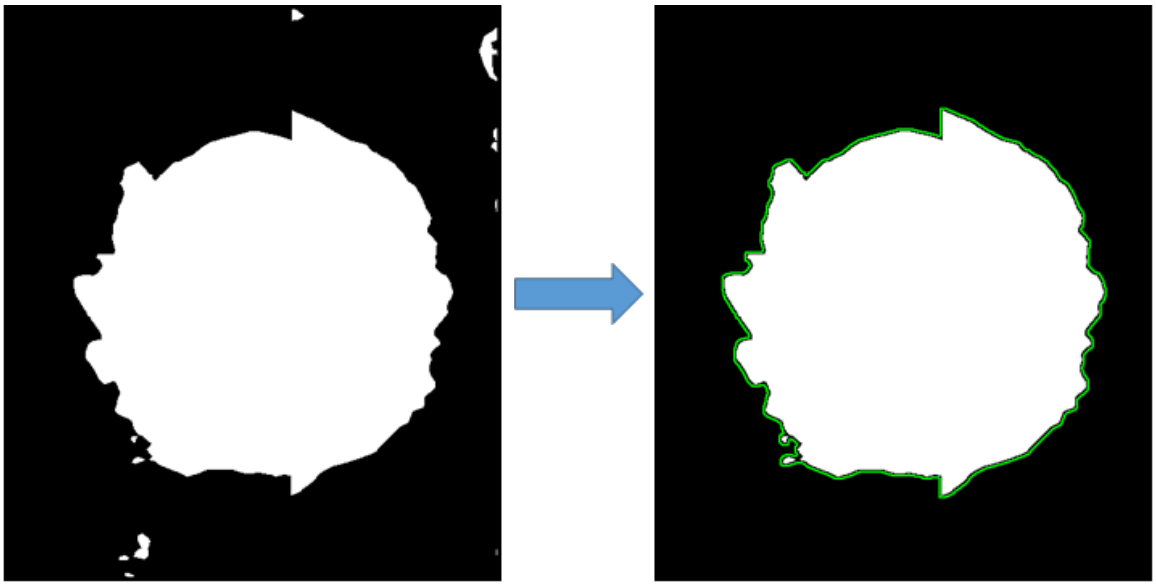

Figure 7.13: Example output of the induration margin identification algorithm, with the selected margin in green pixels.

As shown in Figure 7.13, the resulting image does not contain unconnected misclassifications and depicts the rough margin of the induration by green pixels. An obvious flaw of this algorithm is that it would fail if one of the unconnected clusters is larger than the induration. In this case, the algorithm will identify the unconnected cluster as the induration. However, given that one of the stipulations in capturing the top view image is

that the induration must cover as much of the induration box on the image capture screen as possible, it follows that the induration should be the largest object in the cropped frame.

7.2.5 Step 5: Elliptical approximation of induration

Segmented images of the induration often present with jagged edges. Upon observation, it was found that these edges inaccurately increase or decrease the width of the induration. Such jagged edges emerge because of slight errors in segmentation. As mentioned in the literature review (see Section 2.1), TST induration typically present in a form resembling an ellipse. Fitting an ellipse to the segmented induration is a useful tool for averaging the segmentation error and approximating the edge of the induration.

Elliptical approximation of the edge of the induration is conducted using the EllipseFitter Python package which can be altered to fit an ellipse on a set of binary data points in a 2D matrix representing the segmented image. The approximated ellipse has the same area as the segmented induration, and the ellipse is positioned and orientated such that its area is distributed across the image plane in a similar manner to the area of the segmented induration.

Figure 7.14 shows the result of the ellipse fitting algorithm.

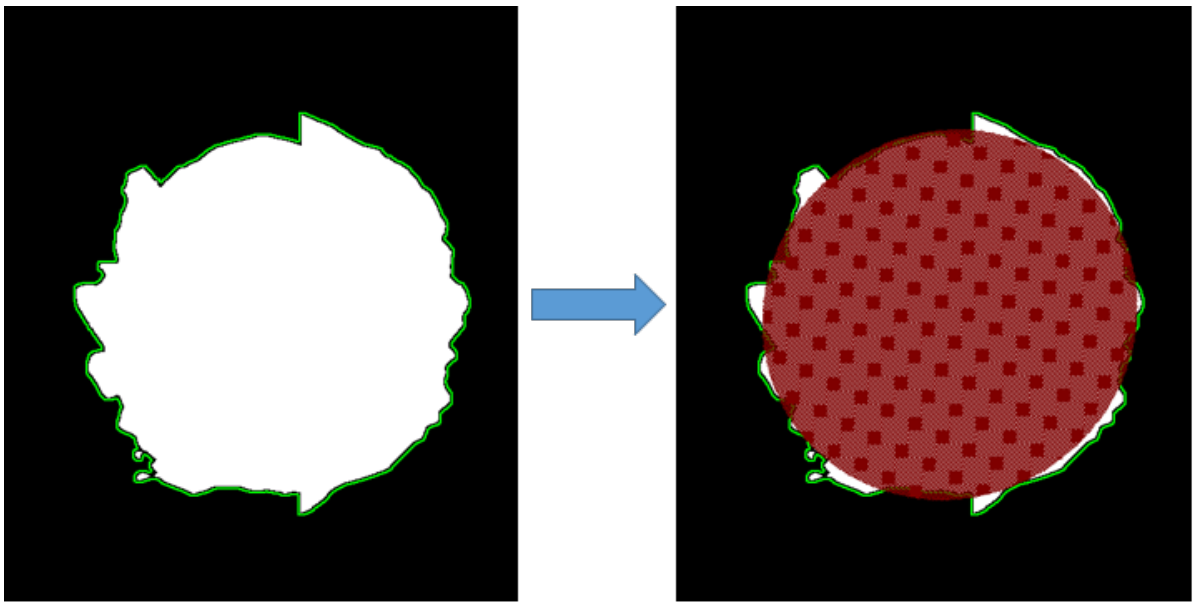

Figure 7.14: Output of ellipse fitting algorithm used to estimate the margins of the induration.

7.3 Induration measurement

7.3.1 Determining the long axis of the arm

The elliptical estimation produced at this point was used to measure the width of the induration. However, as previously mentioned, this width should be measured perpendicular to the long axis of the arm. The long axis of the arm can be defined in two ways: the long axis defined by the image capture screen (capture screen axis), and the long axis defined by the scaling sticker (scaling sticker axis). Figure 7.15 shows examples of these two long axes of the arm on the same depth map image.

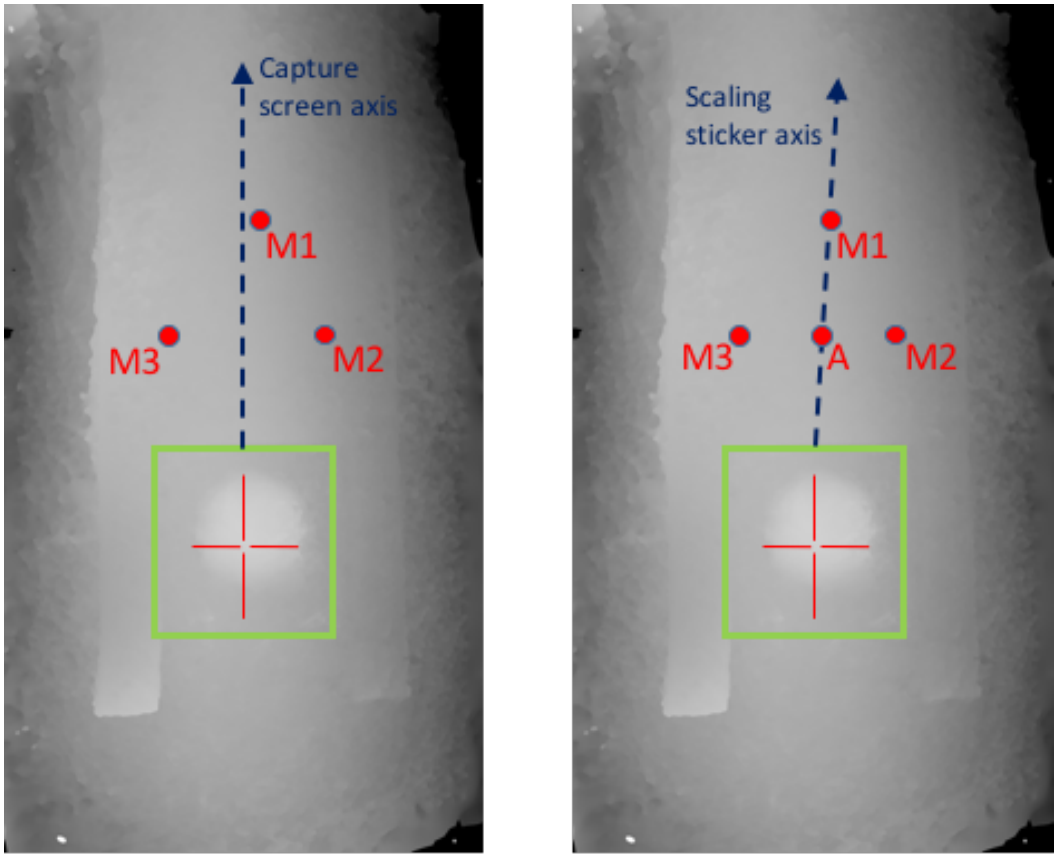

Figure 7.15: Examples of capture screen and sticker long axes of the arm on the same depth map image.

As shown in Figure 7.15, the capture screen axis is easily defined as it runs vertically upwards. It thus follows that, for this case, the width of the induration can be measured perfectly horizontally across the image. The scaling sticker axis, however, requires additional processing to define its direction for measurement purposes. The gradient of the scaling sticker axis is determined (relative to the image pixel Cartesian coordinate system) by

finding point A (the midpoint of M2 and M3) and calculating the gradient between point A and marker M1. The induration is measured along a line perpendicular to this gradient. In the experiments to test the accuracy and functionality of the developed screening tool (Chapter 8), the induration is measured using both axes and their results compared to the width of the induration as measured by a clinician. This provides an indication of which method of defining the long axis of the arm is more suitable for measuring the width of the induration using the developed screening tool.

7.3.2 Measurement and recording of results

The final phase of the measurement process is concerned with measuring the width of the induration and recording the results and intermediate results of the process. The diagnostic width of the induration was defined as the largest width of the elliptical approximation of the induration perpendicular to the long axis of the arm. It follows that two such measurements can be made; according to the capture screen axis and according to the scaling sticker axis. Figure 7.16 shows an example of an induration with its width marked according to both the abovementioned axes.

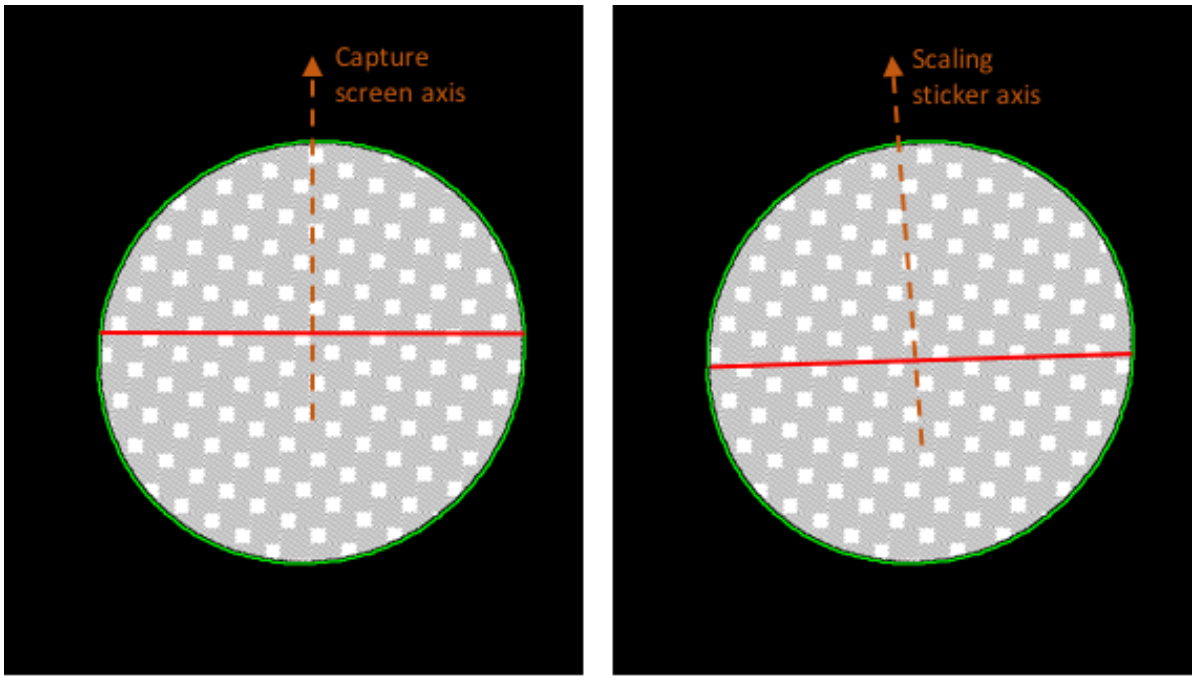

Figure 7.16: Induration measurements according to capture screen and scaling sticker axis.

The red lines (Figure 7.16) represent the largest width of the induration in relation to the capture screen axis and scaling sticker axis respectively. These lines can be measured by the

number of pixels along their breadth. Using the scale of the depth map (in mm/pixel) and the number of pixels in the red lines, a real-world measurement of the induration is determined in millimetres.

Once the measurement process is complete, all images showing measurements and intermediate image processing results (cropped and histogram equalised induration images) are saved to the patient-specific Depth Map folder. An excel spreadsheet is constructed, and saved in the Output Data folder, which contains the following data in each of its sheets:

- Image Data (sheet 1): Pitch and roll values in degrees for each image used in the reconstruction.
- Marker Data (sheet 2): Marker pixel coordinates (in pixels) of the 3 markers, the marker to marker distance (in pixels), the gradient of scaling sticker long axis line and the scale of depth map image (in mm/pixel).
- Questionnaire Data (sheet 3): Questionnaire questions and answers as received from the mobile application and minimum induration size for positive screening result (derived from the answers to the questionnaire questions).
- Results (sheet 4): Pixel and millimetre measurement of induration according to both arm long axes, screening result (positive or negative) and a comparison of induration measurement by the algorithm to the clinician's measurement of the induration.

8. SIMULATED CLINICAL TESTING AND DISCUSSION

This chapter presents the methodology, results and discussion of two experiments used to test the accuracy of the developed evaluation system in ideal conditions and simulated real-world conditions, respectively. In addition to assessing the accuracy of the evaluation system, experiment 2 assesses the usability of the application and aspects of the physical nature of the induration which influence the final measurement result.

8.1 Experiment 1: Accuracy experiment

8.1.1 Background

An experiment to test the accuracy of the TST evaluation system was conducted with the specific aims of:

1. Quantifying the accuracy of the developed TB screening system.
2. Comparing the accuracy of the result (size of induration) for 3, 5 and 7 image inputs in the photogrammetric 3D reconstruction process.
3. Determining the minimum number of images (from the set of 7 images) needed for accurate measurement of the TST induration (within limits specified in the system design criteria – Chapter 3).

To quantify the accuracy of the evaluation system, the experiment was conducted on an induration of known width. This was achieved by modelling an arm surface and induration in a 3D modelling software and 3D printing the model for use in this experiment. The induration on the 3D printed model had a height and width of 4mm and 19mm respectively. The selected height and width of the induration were chosen to represent a condition where the induration is well-elevated and large. The following apparatus (depicted in Figure 8.1) was used for the experiment:

1. HTC One M8 smartphone loaded with TB screening application.
2. Orange 3D printed arm surface with a simulated induration.
3. Scaling sticker.
4. White paper sheet.

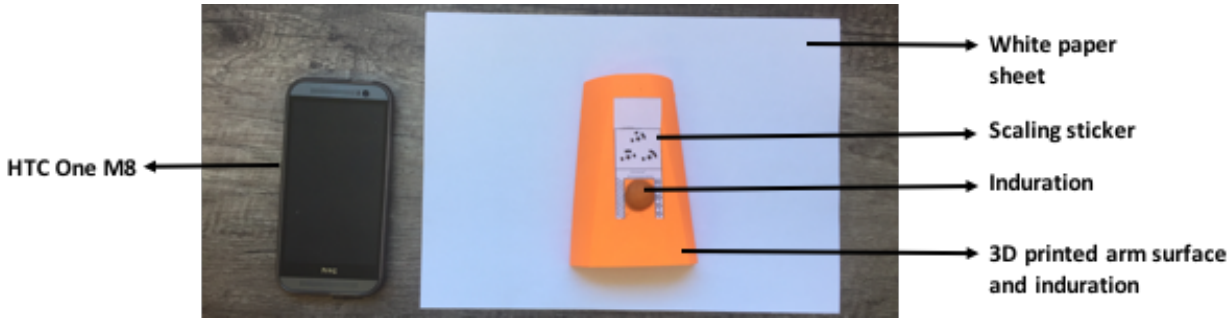

Figure 8.1: Experiment apparatus.

8.1.2 Methods

The experiment was conducted in a well-lit room on a flat table. The 3D printed arm surface with simulated induration was placed on a white sheet of paper on top of the flat table as shown in Figure 8.1. Using the developed application, and following the image acquisition protocol, 7 images of the induration were captured. The images were captured at predetermined pitch and roll angles as presented in Figure 8.2 below.

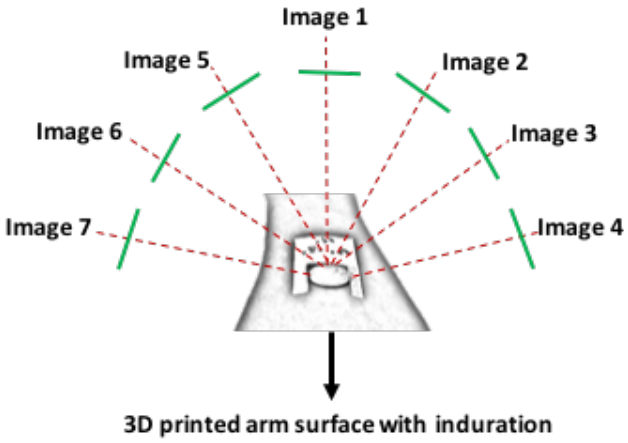

Image	Pitch (Degrees)	Roll (Degrees)
Image 1	0	0
Image 2	0	25
Image 3	0	39
Image 4	0	49
Image 5	0	-25
Image 6	0	-39
Image 7	0	-49

Figure 8.2: General orientation of the captured images (left) and their predetermined pitch and roll values (right).

Ten sets of 7 images were captured (at the specified pitch and roll angles shown in Figure 8.2 – left) using the application. The image sets were uploaded onto a desktop computer and each of the 10 experiment image sets was run through the developed processing pipeline, which produced a 3D reconstruction of the simulated induration and measured the induration. In the next step, two images (image 4 and image 7) were removed from each image set and these modified image sets were run through the developed processing pipeline (for a 5-image reconstruction). In the third step, a further two images (image 3 and

image 6) were removed from the 5 image sets and processed using the developed processing pipeline (for a 3-image reconstruction). Finally, the measurement mean, standard deviation and the minimum and maximum measurements for each reconstruction variant (3, 5 and 7 images) were calculated.

8.1.3 Results and discussion

Figure 8.3 presents a comparison of measurement errors between the 3-image, 5-image and 7-image reconstructions.

Figure 8.3: Measurement error comparison between 3, 5 and 7 image reconstructions.

Table 8.1 shows the measurement mean, standard deviation (SD) and the minimum and maximum measurements for each reconstruction variation (3, 5 and 7 images); the actual measurement was 19 mm.

Table 8.1: Mean, standard deviation, and minimum and maximum measurements for 3 the reconstruction variations in experiment 1.

	7-Image reconstruction	5-Image reconstruction	3-Image reconstruction
Mean measurement (mm)	19.1	18.7	17.1
Measurement SD (mm)	0.3	0.5	1.2
Minimum measurement (mm)	18.7	18.1	16.2
Maximum measurement (mm)	19.7	19.9	19.4

Since each image set was captured in the same lighting conditions and at consistent specified pitch and roll angles, measurement results could be compared without the influence of environmental conditions. It was found that the 7-image reconstruction had the lowest measurement standard deviation (0.3 mm) followed by the 5 and 3 image reconstructions with standard deviations of 0.5 mm and 1.2 mm, respectively. The 7-image reconstruction resulted in a mean measurement (19.1 mm) which was closest to the true width of the induration (19 mm); the 5 and 3 image reconstructions had mean measurement values of 18.7 mm and 17.1 mm, respectively. This result was anticipated as more images (and thus more information about the induration and arm surface) in the photogrammetric 3D reconstruction was expected to result in a better 3D reconstruction and a more accurate depth map; thus resulting in a more accurate measurement of the induration.

Although the 7-image reconstruction results had the least measurement error for 8 of the 10 image sets, the 5-image reconstruction resulted in a lower measurement error for image sets 8 and 9 (Figure 8.3). This suggests that, in these image sets, sufficient depth information about the arm surface was captured in the 5 central images (refer to Figure 8.2 – left) and that the two edge images (image 4 and image 7) may have contributed incorrect depth information which distorted the 3D reconstruction. Upon visual inspection of these image sets it was found that 5 central images were more in focus and clearer than the two edge images. Furthermore, as shown in Figure 8.3, image sets 5 and 8 resulted in low measurement errors (below 0.4 mm) for all 3 reconstruction variations. It was found that the images of these image sets were visually clearer and more in focus as compared to the other image sets. This suggests that the sharpness and the degree to which the images are in focus play a role in the induration measurement result. This result was expected as the SFM photogrammetric process used for 3D reconstruction identifies more distinct features in clearer images; resulting in more tie-points and thus a better 3D reconstruction of the arm surface and induration. The variation in sharpness and focus across the 10 image sets could be due to two reasons: the stability of the operator's hand at the point of capture and failure of the mobile application's autofocus function to focus sufficiently.

The experimental conditions (outlined below) of this test simulated the best-case operational scenario for the developed application. In particular, the images used in this experiment were captured by an operator other than the subject (not captured by someone imaging their own arm). In real-world use, the application may be operated by the patient. In such a case, the quality of the captured images (sharpness and focus) may be compromised. Furthermore, the 3D printed induration had well-defined margins (the edges of the induration were clear and visually identifiable) and well-raised off the arm surface (4 mm height). A real-world TST induration, however, may be shallow and present with edges that are not well defined. The following experiment simulates real-world conditions.

8.2 Experiment 2: Skin induration experiment

8.2.1 Ethics approval

Participants for this experiment were recruited from the postgraduate student cohort in the Division of Biomedical Engineering at the University of Cape Town. Ethics clearance for this experiment was obtained from the human research ethics committee of the University of Cape Town (HREC REF: 890/2016). Prior to conducting the experiment, the participants were informed of the purpose and procedure for this experiment. The participants signed a consent form prior to the experiment.

8.2.2 Background

An experiment to test the evaluation system in real-world conditions was conducted with the aim of assessing the performance of the system when:

1. The patient images their own induration.
2. Imaging indurations of various heights and diameters.
3. Imaging indurations without well-defined visually identifiable margins.

In addition, the experiment also aimed to:

4. Compare measurements made relative to the sticker and capture screen long axis of the arm, to manual measurements by a clinician.
5. Compare the use of 3, 5 and 7 image reconstructions under real-world conditions.
6. Assess the usability of the developed application.

To simulate real-world conditions a special effects make-up artist (under the guidance of a clinician) was used to create mock indurations on the forearms of experiment participants. The HTC One M8 smartphone was used to capture the images for this experiment. The images were captured on a table covered in white paper.

8.2.3 Methods and materials

The special effects make-up artist constructed indurations on the forearms of ten participants of varying skin tones. Figure 8.4 shows the special effects artist applying an induration to the arm of one of the experiment participants.

Figure 8.4: Special effects artist applying induration to the arm of an experiment participant.

Under the supervision of the clinician, indurations of varying height, shape and margin definition (prominence of edges) were applied to the forearm of the participants. In addition, some indurations were made to look glossy to simulate an infected appearance. Examples of the mock indurations on the participant's arms are shown in Figure 8.5.

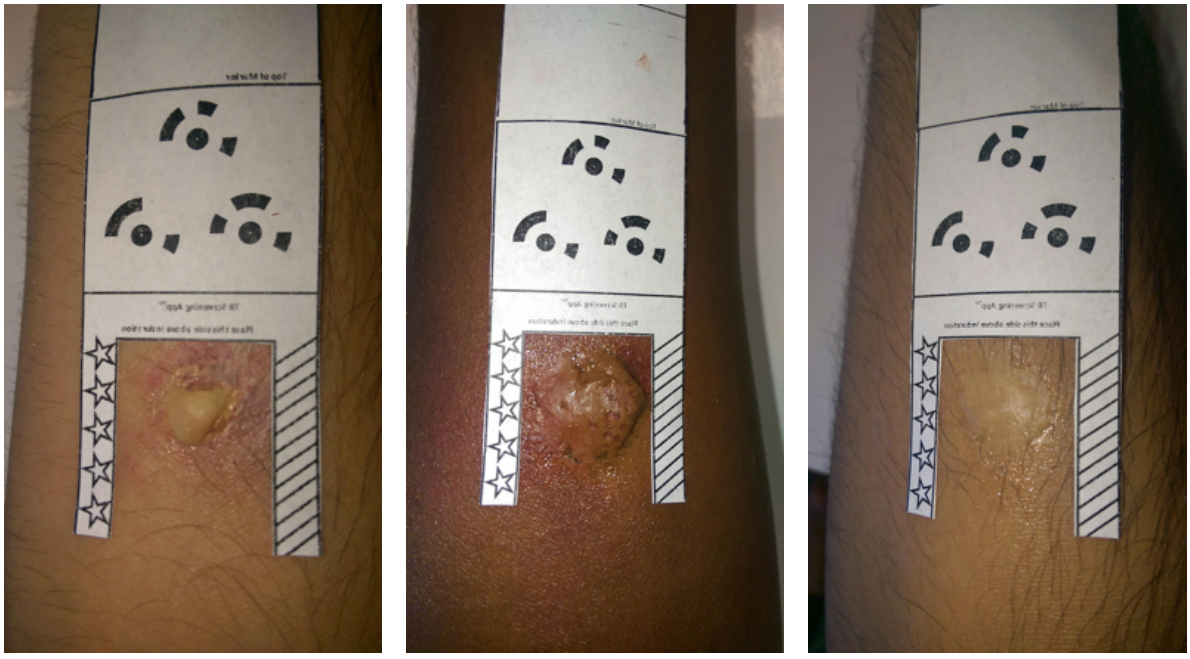

Figure 8.5: Examples of mock indurations on the forearms of experiment participants.

The scaling sticker used in this experiment was modified by removing the line indicating the direction of the long axis of the arm. Figure 8.6 shows the original design of the scaling sticker and the modified scaling sticker used in this experiment.

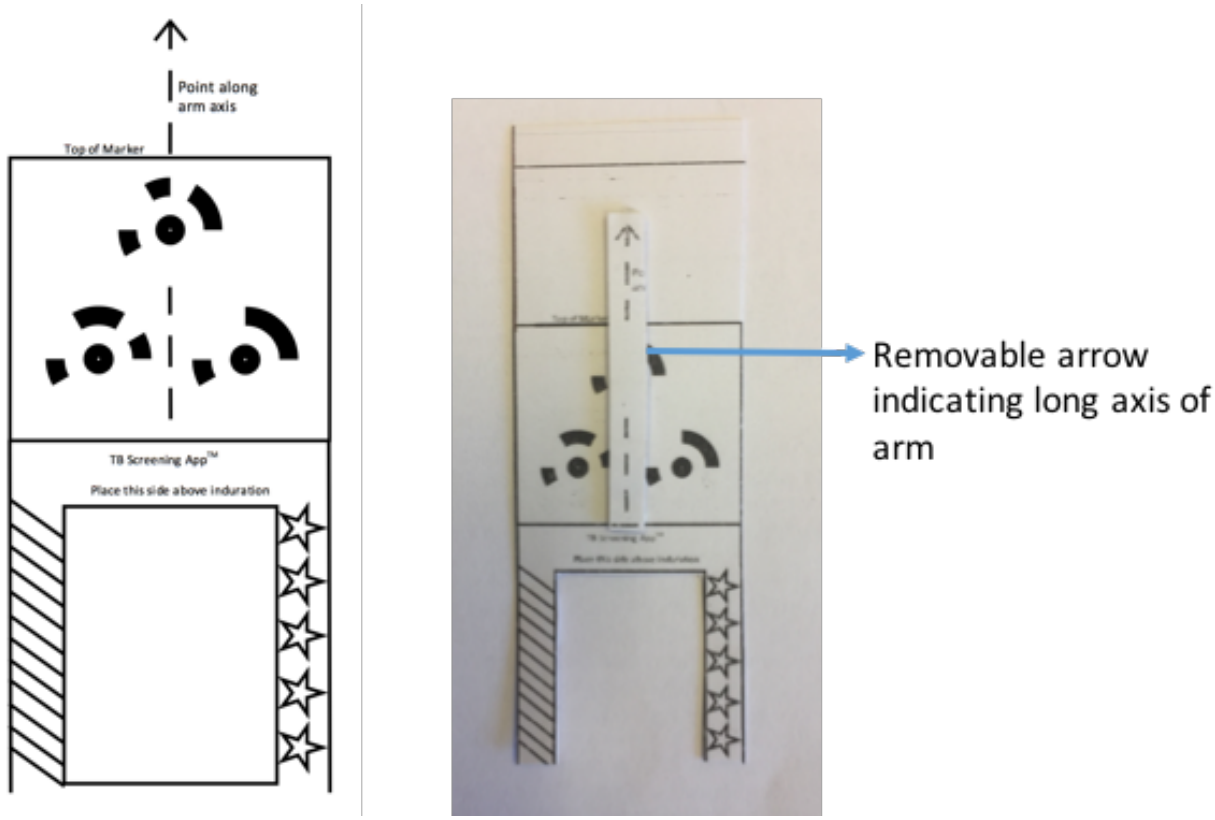

Figure 8.6: Images of original scaling sticker design (left) and modified scaling sticker for use in experiment 2 (right).

As shown in Figure 8.6, the scaling sticker was modified so that the arrow indicating the long axis of the arm can be stuck on the sticker and removed. The arrow was initially stuck on the scaling sticker and each participant was asked to apply the scaling sticker to their forearm such that the arrow points along the long axis of the arm as best as possible. Once the sticker was applied, the arrow was removed from the scaling sticker. In the absence of this modification to the scaling sticker, the participants may have tended to align the image capture screen long axis of the arm with the arrow on the scaling sticker. This would mean that measurement of the induration width in relation to both these axes would be the same. Since an aim of this experiment is to assess which of the arm long axes (capture screen and scaling sticker long axes of the arm) best approximate measurement of the induration made by a clinician, it was desired that the orientation of either axis was not influenced by the other.

Each induration was measured by a clinician using the ruler and pen evaluation method. Following this, each participant was asked to use the developed application to capture the 7 images of the induration according to the image capture protocol. The indurations were removed from the arm after the image capture process and the height of each induration was measured using a Vernier caliper. The image sets were downloaded onto a desktop computer and run through the developed processing pipeline. Following the methodology used in experiment 1, 3-image, 5-image and 7-image reconstructions were produced for each image set. The produced 3D reconstructions and segmented depth maps were visually assessed to verify that the 3D reconstruction was successful and that the segmented depth maps approximate the margins of the induration. Unsuccessful 3D reconstructions and depth map segmentations were noted and the induration measurements (in relation to both long axes of the arm) for the successful image sets were recorded.

Finally, all participants were asked to complete a questionnaire relating to the usability of the developed application. The questionnaire is shown in Appendix B.

8.2.4 Results and discussion

3D reconstruction and depth maps

Table 8.2 shows the physical characteristics of each of the 10 mock indurations and a summary (success or failure) of the segmentation process on the generated depth map for each image set.

Table 8.2: Summary of mock indurations physical characteristics and results of the depth map segmentation.

Image set	Induration characteristics			Good induration segmentation		
	Description of induration	Induration width (clinician measurement)	Induration height	7-image	5-image	3-image
1	Good edge definition Not infected Large width Pronounced height	15 mm	4 mm	✓	✓	✓
2	Good edge definition Not infected Moderate width Pronounced height	12 mm	3 mm	✓	✓	✓
3	Moderate edge definition Infected Large width Moderate height	15 mm	2.5 mm	✓	✓	✓
4	Moderate edge definition Infected Moderate width Shallow height	11 mm	1.5 mm	✓	✗	✗
5	Moderate edge definition Not infected Moderate width Shallow height	8 mm	1.5 mm	✓	✗	✗
6	Moderate edge definition Not infected Small width Shallow height	7 mm	1.2 mm	✓	✓	✗
7	Bad edge definition Not infected Small width Shallow height	3 mm	0.5 mm	✗	✗	✗
8	Moderate edge definition Infected Small width Shallow height	7 mm	1 mm	✓	✗	✗
9	Moderate edge definition Not infected Small width Shallow height	6 mm	1 mm	✗	✗	✗
10	Bad edge definition Not infected Small width Shallow height	7 mm	0.3 mm	✗	✗	✗

Legend: ✓ Success ✗ Failure

The 3D reconstruction in APS was successful for all the image sets across all the reconstruction configurations (3, 5 and 7 image reconstructions). This was assessed by visually confirming that the surface of the arm, scaling sticker and induration was successfully reconstructed in APS. However, in some cases, good induration segmentations were not produced from the segmented depth map images. Figure 8.7 shows examples of

some induration segmentation failures alongside the unsegmented depth map of the induration from the experiment data set.

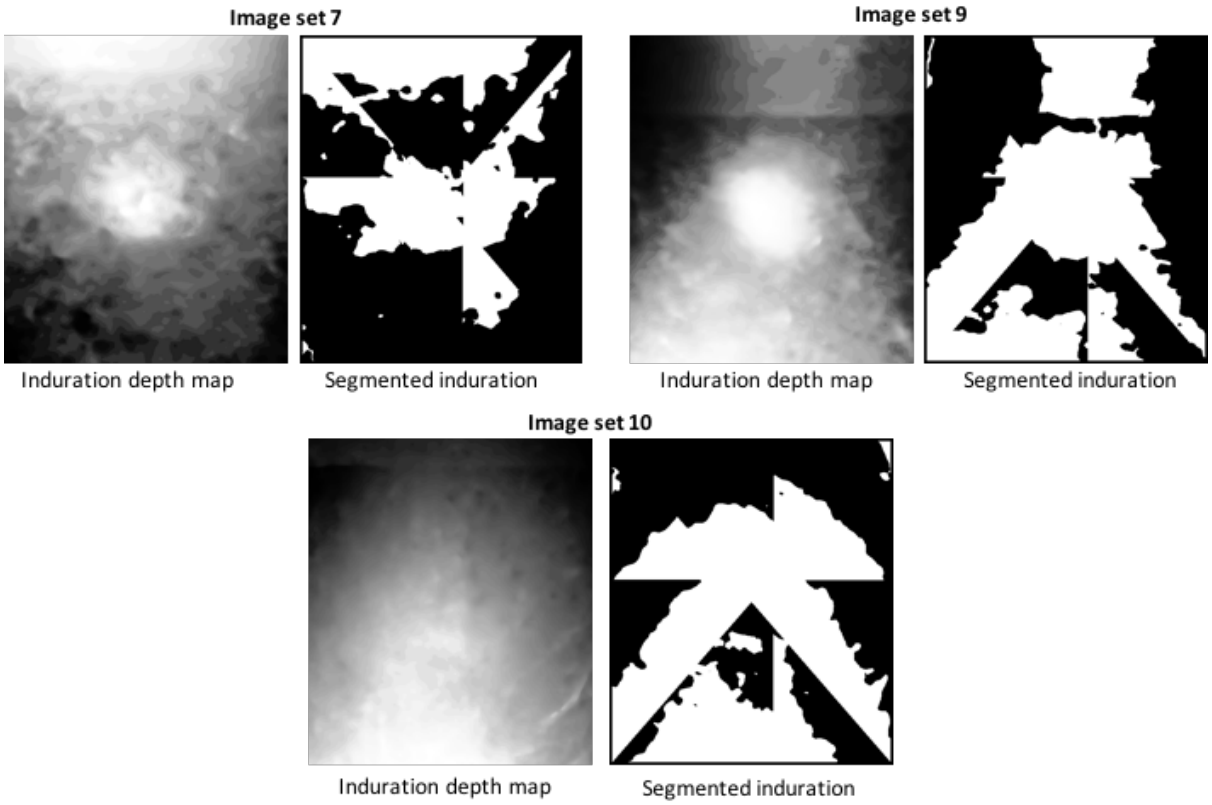

Figure 8.7: Induration depth map and segmentation failure for image sets 7, 9 and 10.

Sixteen of the 30 reconstructions resulted in segmentation errors similar to the ones shown in Figure 8.7. This result is not consistent with the results presented in experiment 1, where all image sets resulted in correctly segmented depth map images of the induration. Upon comparison of the images used in experiment 1 to the images used in this experiment, it was found that the images used in this experiment were generally blurrier, which was mainly attributed to the stability of the device at the point of capture of the images. Since the participants were imaging their own arms in this experiment, the loss of sharpness of the images could be due to the fact that the operator (experiment participant) was focused on multiple tasks during the image capture process. The participant had to ensure that the arm with the induration was still, the image capture specifications were met, and the mobile device was stable, at the point of capture. This problem can be mitigated by requiring that an external operator capture the images for the patient. This may result in an increase in stability at the point of capture and thus increase the sharpness of the captured images.

The errors in the segmentation process (Figure 8.7) were more prominent for the 3 and 5 image reconstructions than for the 7 image reconstructions; with 7, 6 and 3 failed segmentations, respectively. It was expected that the number of tie points and dense cloud points would decrease with a reduced number of input images. This resulted in depth maps with low contrast between the induration's raised surface and the surrounding skin. This is of significance when imaging low height indurations where low contrast between the induration surface and surrounding skin is expected in the depth map regardless of the number of images used in the reconstruction process. Because of the lack of sufficient contrast between the pixels representing skin and induration, the segmentation fails. This problem could be rectified by using larger image sets to increase the quality of the 3D reconstruction and the produced depth map.

However, another possible solution to this problem, which can be implemented on image sets of the same size as the ones used in this experiment, is improving the segmentation algorithm. As seen in the depth maps of image sets 7 and 9 (Figure 8.7), the induration is visually identifiable near the centre of the depth map image, however, the lighter patches (below and above the induration) on the depth map negatively impact the segmentation. Combining a machine learning algorithm, which learns how to correctly identify the induration (by learning common shapes and locations of the induration in the depth map), with the segmentation method may address this problem.

It was found that signs of infection on the mock indurations did not influence the performance of the 3D reconstruction process and the resulting depth maps in this experiment. However, it should be noted that only two such indurations were simulated in this experiment and that the variation of all possible signs of infection were not represented. The factors that had the greatest impact on the 3D reconstruction in APS and the produced depth maps of the induration were the definition of the induration's edges and the height of the induration. Figure 8.8 shows examples of two indurations from the image sets used in this experiment. The left induration (Figure 8.8) represents a well-elevated (high) induration with good edge definition while the induration on the right represents a shallow induration with poor edge definition.

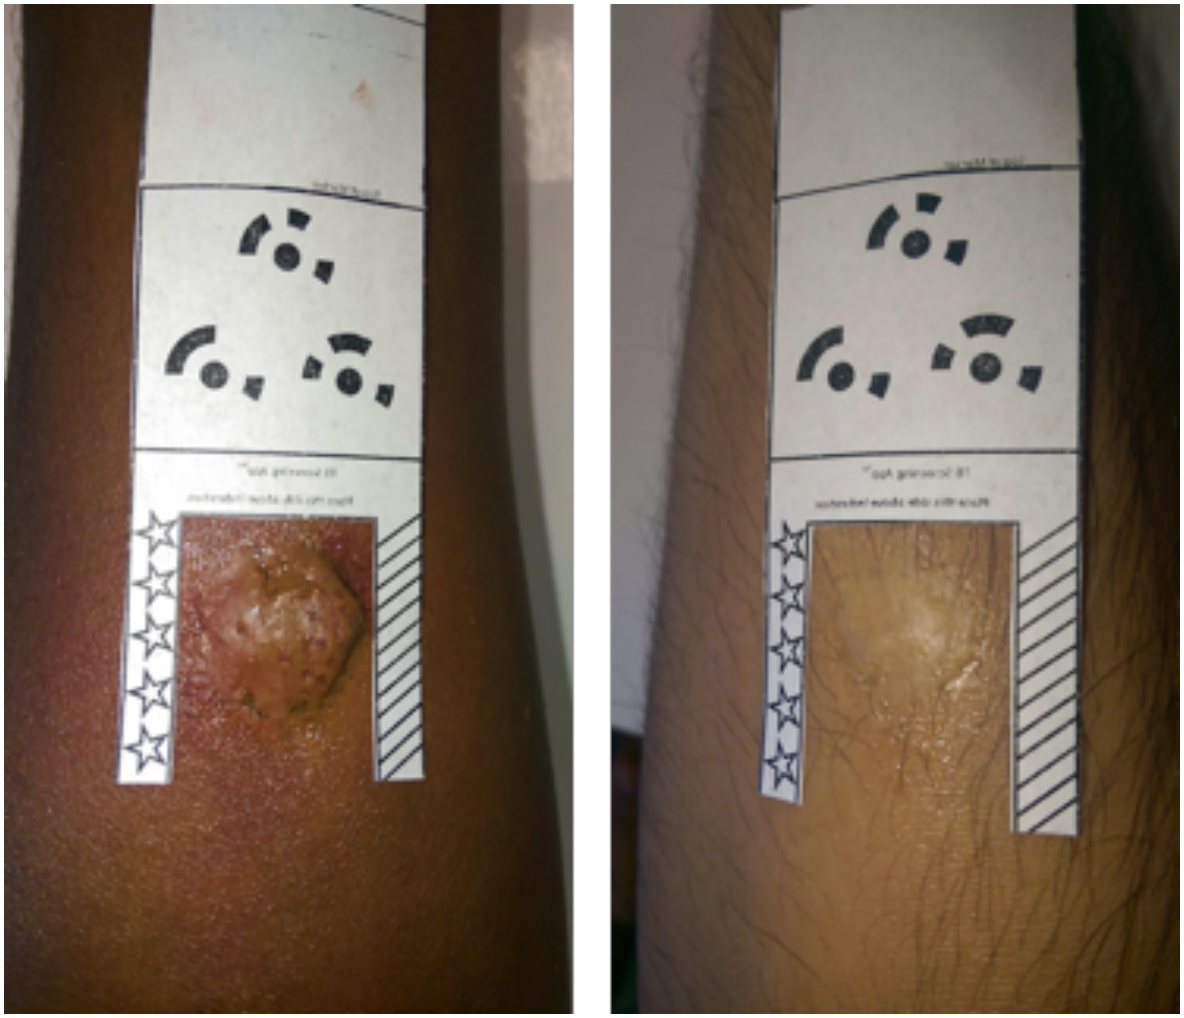

Figure 8.8: Examples of well elevated induration with well-defined edges (left) and shallow induration with poor edge definition (right).

The reason these factors (induration height and edge definition) had a significant influence on the reconstruction process (and thus the produced depth map of the induration) was probably due to the low resolution of the HTC One smartphone camera used to image the indurations. Low-resolution cameras impart less information about the depth change and surface detail of the imaged object to the photogrammetric 3D reconstruction process. Due to the low pixel density in the images, much of the depth detail captured in the images is aggregated. Higher resolution cameras would represent more layers of depth change of the imaged object and thus are better suited for representing small depth changes and minute details of the object being imaged. Defined edges may not play such a significant role when using higher resolution images in the reconstruction process as sufficient depth information will be available in such image sets. An illustrative quantitative analysis of the effect of these two factors on the final measurement of the induration is presented below.

Measurement of induration and suitable long axis of arm assessment

Figure 8.9 shows a graph of the measurement error of the successfully segmented indurations for the 3 reconstruction variations (3, 5 and 7 image reconstructions) measured relative to both the sticker and capture screen long axes of the arm.

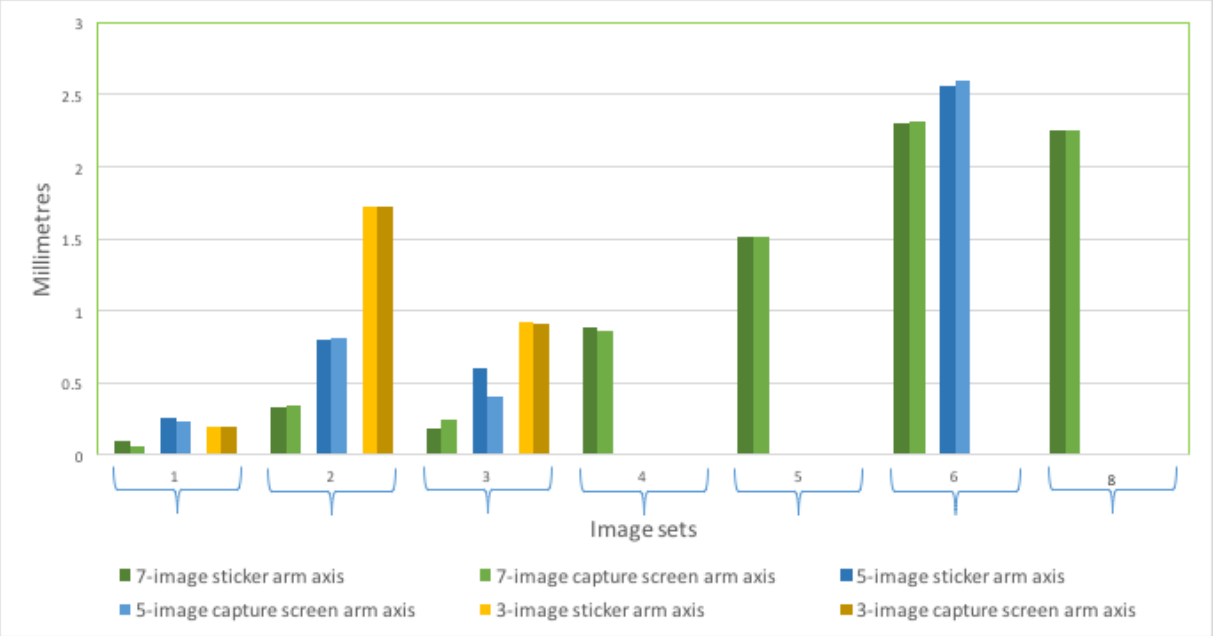

Figure 8.9: Measurement error for 3, 5 and 7 image reconstructions of the indurations measured relative to the sticker and capture screen long axes of the arm in experiment 2.

The missing data points in Figure 8.9 represent the data sets where the depth map segmentation of the induration failed (such as the examples shown in Figure 8.7) and thus the induration was not measured. As seen in Figure 8.9, there is no significant difference between measuring the width of the induration relative to the long axis of the arm derived from the scaling sticker and the long axis of the arm derived from the image capture screen; with the difference between these measurements ranging in 0.00 mm and 0.09 mm. This result suggests that the participants of this experiment perceived the orientation of the long axis of the arm to be the same when applying the scaling sticker and when capturing the top view image. In addition, similar to the results of experiment 1, measurement accuracy increased with increased number of images in the reconstruction process. Table 8.3 shows the number of successful measurements, measurement error mean and standard deviation (SD) for each reconstruction configuration.

Table 8.3: Number of successful measurements, measurement error mean and standard deviation for the 3 reconstruction variations in experiment 2.

	7-Image Reconstruction	5-Image Reconstruction	3-Image Reconstruction
Number of successful measurements	7	4	3
Mean measurement error (mm)	1.1	1.0	0.9
Measurement error SD (mm)	1.0	1.1	0.8

Effect of induration edge definition and induration height on measurement result

As noted in previous sections, the definition of the edges and the height of the induration impact the quality of the produced 3D reconstruction and thus also the quality of the resulting depth map, and ultimately the accuracy of the measurement. To quantify the effect of these two factors on the final measurement of the induration, each induration was allocated an 'edge definition rating'. The rating is a number between 1 and 10, where 1 represented – "poor edge definition", and 10 represented – "good edge definition". By multiplying the edge definition rating with the height of the respective induration, a 'height-edge definition factor' was calculated. Figure 8.10 shows a graph of the induration measurement error vs the height-edge definition factor for the 7-image reconstruction configuration.

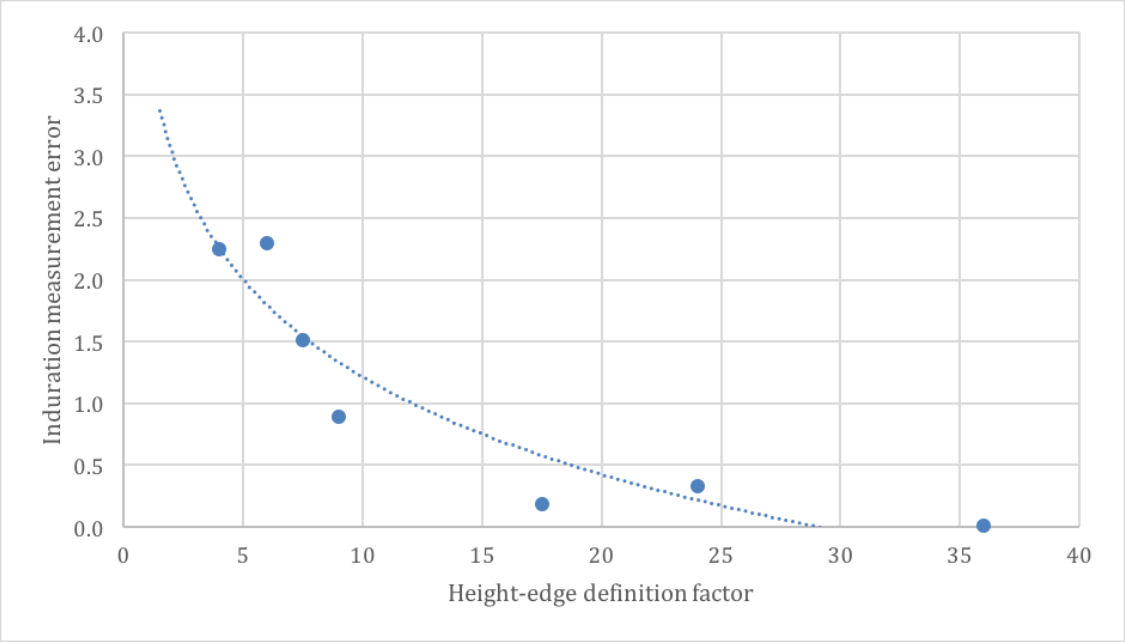

Figure 8.10: Induration measurement error vs height-edge definition factor for 7-image reconstruction configuration.

The measurement data for the 3 and 5 image reconstructions was excluded from the analysis presented in Figure 8.10. The reason for this is that the 3 and 5 image reconstruction configurations had 3 and 4 induration measurement data points respectively and thus there was not sufficient information to draw a trend line for the data. However, upon visual inspection of the data, it was found that the measurement data for the 3 and 5 image reconstructions follow a similar trend to that presented in Figure 8.10.

Figure 8.10 shows that the relationship between the measurement error of the 7-image reconstruction and the height-edge factor of the measured induration is best approximated by an exponential trend line with the measurement error decreasing with increasing height-edge definition factor (higher indurations with more defined edges). As outlined in the section above, less edge definition and lower induration heights are expected to decrease the accuracy the 3D reconstructions of the induration and thus decrease the accuracy of the generated depth maps. With decreasing depth map accuracy, the induration segmentation algorithm resulted in more incorrectly segmented pixels when segmenting the induration from the surrounding skin and hence the higher recorded measurement error.

Usability of developed mobile application

Using the questionnaire shown in Appendix B, the participants were asked to rate the usability of the application according to 10 statements. In the questionnaire the statements were rated on a scale of 1 to 5 where 1 and 5 represents “strongly disagree” and “strongly agree” respectively. The questions in the evaluation form are a mixture of positively and negatively phrased statements. As such, for analysis, the responses to the positively and negatively phased questions were grouped separately. The average response to each question was calculated by summing the scores from 1 to 5 for each question and dividing by the number of responses for the respective questions. The average response value can be represented as a “level of agreement” measure for each question by calculating the average response as a percentage of the upper scale value of 5. Table 8.4 shows the participants’ “level of agreement” with both the positively and negatively phased statements.

Table 8.4: Participants' level of agreement with questions from the TB screening application evaluation form.

Questionnaire statements	Level of agreement
Positively phased questions	
1. I think I would be able to use this application frequently.	78%
2. I thought this mobile app was easy to use.	76%
3. I found the various functions of this mobile app were well integrated.	86%
4. I would imagine that most people would learn to use this mobile application quickly.	70%
5. I felt very confident using this mobile application.	82%
Negatively phased questions	
6. I found this application unnecessarily complex.	36%
7. I think I would need assistance to be able to use this mobile app.	60%
8. I thought there was too much inconsistency in this mobile application.	28%
9. I found this mobile application very cumbersome/awkward to use.	48%
10. I needed to learn a lot of things before I could get going with this mobile application.	48%

The participants of this experiment were well-acquainted with using a mobile phone, fluent in English and familiar with technical terms such as "pitch" and "roll". Most participants had a high level of agreement with the positively phased questions (Table 8.4) with question 4 showing the lowest level of agreement. The general comments concerning this question was that the instruction pages were long and that there was a lot to remember about the imaging protocol. This would make it difficult for people to follow the instructions accurately. This comment aligns with the responses to questions 7 and 10, where participants pointed out that the instructions could possibly be better explained verbally.

Participants also felt that using a cell phone camera is a learnt skill and that people may need to time to adapt to using a mobile phone camera without assistance. The participants noted that, even though they were familiar with using a mobile phone camera, keeping the arm with the induration stable while fulfilling the image acquisition protocol requirements was difficult. These comments correlate with the responses to question 9, where some participants felt that capturing images of their induration with one hand was cumbersome.

Some of the participants felt that the developed application may be unnecessarily complex (question 6) because the user must return to the image capture dashboard after capturing every image. These participants suggested that the user be directly prompted to capture each image directly after the other. This approach, however, would be difficult to implement as the participant would not have time to familiarise themselves with the image specifications for the image before being directed to the capture screen. A possible solution to this problem is replacing the image capture protocol with a video capture protocol.

The participants also pointed out some positive aspects of the application. They noted that the application was visually attractive and well-structured for intuitive use. They also noted that the “tilt direction arrow”, visual cues for indicating the correct pitch and roll values, and guide boxes on the capture screen, were very helpful during the image capture process.

9. General Discussion, Conclusion and Suggestions for Future Work

9.1 Optimal number of images required for complete assessment

Experiment 1 showed that the 7-image reconstruction resulted in the highest accuracy measurement of the induration as compared to the 3 and 5 image reconstructions. As compared with the literature, the 7-image and 5-image reconstructions had lower measurement standard deviations than current methods of measuring the TST induration using a ruler and pen (Pouchot et al., 1997); with 0.3 mm and 0.5 mm measurement standard deviation respectively as compared to 1.1 mm from the literature. However, this experiment was conducted on a single 3D printed induration as opposed to real TST skin indurations used by Pouchot et al. (1997).

Experiment 2 was conducted on mock skin indurations and the results showed that the mean measurement error and standard deviation across the 3 reconstruction configuration (3, 5 and 7 image reconstructions) were similar to or lower than the measurement standard deviation quoted in the literature (Pouchot et al., 1997). However, the calculation of the mean measurement error and standard deviation was distorted by the number of data points (measured induration widths) for each of the reconstruction configurations; with the 3 and 5 image reconstructions resulting in 3 and 4 successful measurements as compared to 7 successful measurements for the 7-image reconstruction. The 3 and 5 image reconstructions failed to produce accurate depth maps for indurations with the lowest height-edge definition factor. This resulted in failed segmentation of the induration margins and an inability to measure the induration width.

In real-world use of the application, the images captured using the application are sent to a central processing unit. It is desired that the cost (internet data transfer cost) of transferring these images is minimised. This can be achieved by decreasing the number of images required for the reconstruction and assessment process. However, the experiments suggest that the higher number of images retained result in a higher measurement success rate.

The cost of transferring additional images to the central processing is directly related to the size (in megabytes) of each additional image. Each image captured using the HTC One M8 smartphone was 0.9 to 1 MB in size. When comparing the 5-image reconstruction to the 7-image reconstruction, the addition of 2 MB in data transfer costs resulted in a 30 percentage point (from 40% to 70%) increase in measurement success rate under the conditions simulated in experiment 2. The size of additional images and measurement success rate (across the 3 reconstruction variations) may differ when using other smartphones in conjunction with the developed application.

Considering the additional data transfer costs, assessment success rate and the measurement accuracy of using 3, 5 and 7 images in the reconstruction process, it was found that the 7-image protocol was best suited for real-world conditions similar to those simulated in experiment 2.

It is suggested that the measurement success rate and measurement accuracy for reconstructions using 3, 5 and 7 images captured using the developed application be compared for smartphones other than the HTC One M8.

9.2 3D reconstruction, induration segmentation and measurement accuracy

The sharpness (focus) of the captured images plays a role in 3D reconstruction quality which in-turn affects the accuracy of the generated depth map and measurement of the induration. The effect of the extent to which the images are in focus (blurriness of images) was seen in both experiment 1 and 2; where blurrier image sets resulted in poor depth maps from the 3D reconstruction and therefore less accurate induration measurements. In this project, the autofocus feature of the smartphone was used to ensure that the captured images were in focus. Although including the autofocus feature removed the need for the operator to manually focus images, some of the captured images may have been out of focus because sufficient time was not given for the smartphone to focus before the image was captured. The ability of a smartphone to autofocus is dependent on the smartphone being used. Factors such as focal length, quality of lens components and processor of the smartphone affect the time taken to autofocus and the sharpness of the image.

Although it was established that image sharpness influences the final measurement of the induration, the effect of image sharpness on the 3D reconstruction, depth map generation and measurement of the induration was not quantified. Such an assessment of the effect of image sharpness on the entire developed screening process would be valuable and can be conducted by adding artificial blurring to the image set to simulate out-of-focus images.

Alongside the sharpness of the images in the image set, the resolution of captured images may influence the final measurement of the induration. Since features on the skin (induration and edges of induration) are relatively small, APS's ability to identify and accurately represent the depth change of these features is dependent on level of detail captured in the image. The detail captured in images increases with increasing pixel density (resolution of images). In the experiments conducted, the image resolution was kept constant (use of a single smartphone in experiments) and the effect of induration height and edge definition was assessed. It was found that the measurement accuracy worsens exponentially for decreasing induration height and less defined induration edges. However, this trend may differ for higher resolution images. Higher resolution images may capture sufficient surface detail regardless of the physical features of the induration; and thus, other factors such as lighting may play a more significant role in the final measurement of the induration.

The induration segmentation method failed to correctly segment the indurations for all the depth maps produced from the 3D reconstructions in experiment 2. The segmentation method could be improved by using machine learning techniques which could incorporate information about the location and shape of the induration in the segmentation process. The developed image capture protocol specifies that the induration must be in the centre of the induration guide box. Furthermore, TST induration are commonly found to be elliptical. Training a machine learning algorithm to pick up such commonalities in the depth maps could improve the performance of the segmentation algorithm. However, to implement this solution, a sufficiently large data set of induration depth maps would be needed to train such a machine learning algorithm.

9.3 Image acquisition protocol and application user interface

The designed image capture protocol was seen to have an effect on the quality of images captured, with images in experiment 2 being blurrier, possibility due to the lack of stability at point of capture when the subject was imaging their own induration. This hypothesis was substantiated by the participant feedback on the usability of the application where they expressed that the capture process was awkward and cumbersome to complete. A possible solution to this problem is specifying that the images be captured by an external operator instead of the patient themselves.

The participants of experiment 2 did not represent the population of a resource-limited setting. Although these participants were well-acquainted with using smartphone cameras and fluent in English, the participants found the image capture protocol complex and difficult to execute. These problems are expected to be amplified should the application be used by a population in a resource-limited setting where English may not be the first language and capturing images using a smartphone is not common. The issue of language barrier and understanding of the image capture protocol could be addressed through translating the protocol instructions into a language which is well understood by the user. Furthermore, the instruction pages for the image capture protocol could be compiled into an instruction video which the user could watch before attempting the capture process. The problem of unfamiliarity with operating smartphones could be mitigated by using a trained operator to capture the images for people unfamiliar with smartphones.

Participants of experiment 2 reported that the image capture protocol was unnecessarily complex. It was suggested that the image capture protocol be replaced with a single video capture protocol. To implement this change, the patient would be required to capture a video of the induration which covers the 100° range around the induration originally covered by the 7-image capture protocol. An illustration of this is shown in Figure 9.1.

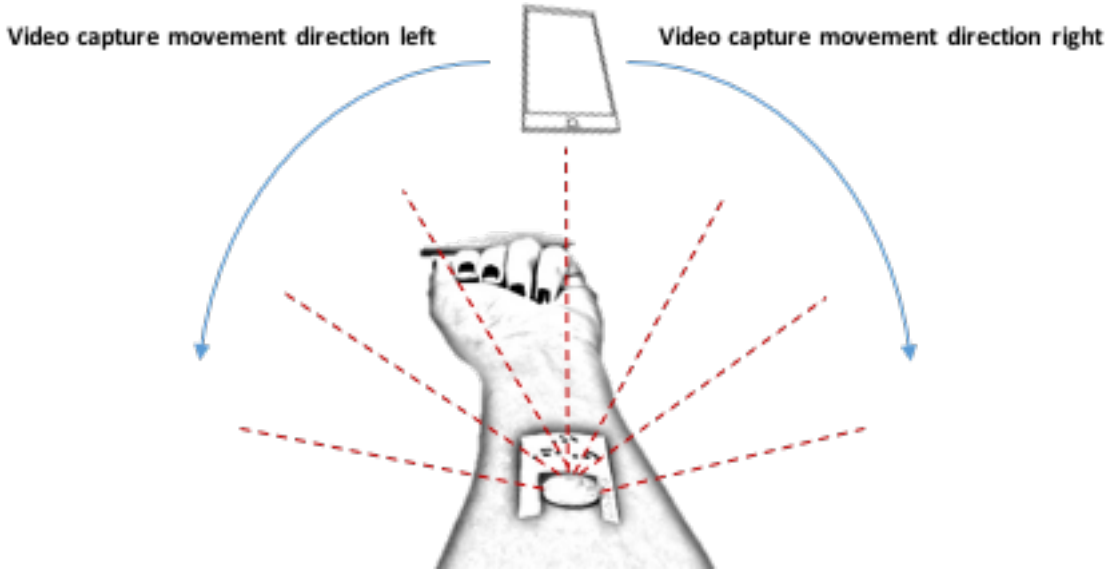

Figure 9.1: Possible video capture protocol to replace developed image capture protocol.

In the case depicted in Figure 9.1, the user could be guided in capturing a single video of the induration and arm surface. Frames from this video (at specified pitch and roll values) would be extracted from the video and use in the photogrammetric reconstruction. Although this method would make the capture process easier, an effort should be made to ensure the clarity and sharpness of the selected frames of the video, which would be used for the reconstruction process. Since the captured video would pan around the induration to capture the 100° range originally covered by the 7-image protocol, the movement of the video might be a significant source of blurriness in the selected frames.

The processing conducted on the remote processing unit (3D reconstruction, depth map generation and induration measurement) is currently implemented through a Python script. For real-world use, a graphical under interface (GUI) would need to be developed for the processing phase conducted on the remote processing unit. This would allow minimally trained personnel to initiate the screening process using a set of images retrieved from the developed application. Furthermore, the GUI would allow easy access and viewing of the results and intermediate results of the screening process by a clinician.

9.4 Summary of suggestions for future work

The major issues to be addressed in future work are:

- Testing using other smartphones than the model used for this project.

- Revision and optimisation of the depth map segmentation method used in this project.
- Reconsideration of how the images of indurations would be taken in practice.
- Further analysis of the optimal number of images required for analysis of the test results for a single patient.
- Development of a GUI for the system to allow minimally trained personnel to initiate the screening process using a set of images retrieved from the developed application.

10. REFERENCES

ABDOOL KARIM, S. S., NAIDOO, K., GROBLER, A., PADAYATCHI, N., BAXTER, C., GRAY, A. L., GENGIAH, T., GENGIAH, S., NAIDOO, A. & JITHOO, N. 2011. Integration of antiretroviral therapy with tuberculosis treatment. *New England Journal of Medicine,* 365**,** 1492-1501.

AISU, T., RAVIGLIONE, M. C., VAN PRAAG, E., ERIKI, P., NARAIN, J. P., BARUGAHARE, L., TEMBO, G., MCFARLAND, D. & ENGWAU, F. A. 1995. Preventive chemotherapy for HIV-associated tuberculosis in Uganda: an operational assessment at a voluntary counselling and testing centre. *Aids,* 9**,** 267.

BALLABENI, A., APOLLONIO, F., GAIANI, M. & REMONDINO, F. 2015. Advances in image pre-processing to improve automated 3D reconstruction. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences,* 40**,** 315.

BARGAL, S. A., WELLES, A., CHAN, C. R., HOWES, S., SCLAROFF, S., RAGAN, E., JOHNSON, C. & GILL, C. Image-based ear biometric smartphone app for patient identification in field settings. Proc. International Conf. on Computer Vision Theory and Applications, 2015.

BOLLES, R. C., BAKER, H. H. & MARIMONT, D. H. 1987. Epipolar-plane image analysis: An approach to determining structure from motion. *International Journal of Computer Vision,* 1**,** 7-55.

BRESLAUER, D. N., MAAMARI, R. N., SWITZ, N. A., LAM, W. A. & FLETCHER, D. A. 2009. Mobile phone based clinical microscopy for global health applications. *PloS one,* 4**,** e6320.

CAHN, P., PEREZ, H., BEN, G. & OCHOA, C. 2003. Tuberculosis and HIV: a partnership against the most vulnerable. *Journal of the International Association of Physicians in AIDS Care (JIAPAC),* 2**,** 106-123.

CDC. 2013. *Mantoux Tuberculin Skin Test Chart* [Online]. Available: https://www.cdc.gov/tb/publications/posters/images/mantoux_wallchart.pdf [Accessed 8 August 2016].

CHAISSON, R. E., KERULY, J. C., MCAVINUE, S., GALLANT, J. E. & MOORE, R. D. 1996. Effects of an incentive and education program on return rates for PPD test reading in patients with HIV infection. *JAIDS Journal of Acquired Immune Deficiency Syndromes,* 11**,** 455-459.

CHEN, B., FU, Z., OUYANG, H., SUN, Y., GE, C. & HU, J. Fish measurement using Android smart phone: the example of swamp eel. Third International Conference on Photonics and Image in Agriculture Engineering (PIAGENG 2013), 2013. International Society for Optics and Photonics, 876110-5.

DENKINGER, C., GRENIER, J., STRATIS, A., AKKIHAL, A., PANT-PAI, N. & PAI, M. 2013. Mobile health to improve tuberculosis care and control: a call worth making [Review article]. *The International Journal of Tuberculosis and Lung Disease,* 17**,** 719-727.

DESEILLIGNY, M. P. & CLERY, I. Apero, an open source bundle adjusment software for automatic calibration and orientation of set of images. Proceedings of the ISPRS Symposium, 3DARCH11, 2011, 269-276.

FOLARANMI, T. 2014. mHealth in Africa: challenges and opportunities. *Perspectives in public health,* 134**,** 14-15.

GAMADIA, M., KEHTARNAVAZ, N. & ROBERTS-HOFFMAN, K. 2007. Low-light auto-focus enhancement for digital and cell-phone camera image pipelines. *IEEE Transactions on Consumer Electronics,* 53.

GILLWALD, A., STORK, C., CALANDRO, E. & GILLWALD, A. 2013. Internet going mobile: internet access and use in 11 African countries. *info,* 15**,** 34-51.

HAGHDOOST, A. A., AFSHARI, M., BANESHI, M. R., GOUYA, M. M., NASEHI, M. & MOVAHEDNIA, M. 2014. Estimating the Annual Risk of Tuberculosis Infection and Disease in Southeast of Iran Using the Bayesian Mixture Method. *Iranian Red Crescent Medical Journal,* 1-6.

HAM, C., LUCEY, S. & SINGH, S. 2015. Absolute Scale Estimation of 3D Monocular Vision on Smart Devices. *Mobile Cloud Visual Media Computing.* Springer, 329-353.

HOSMAN, L. & FIFE, E. 2012. The potential and limits of mobile phone usage for development in Africa: Innovation and top-down-meets-bottom-up partnering. *The Journal of Community Informatics,* 8.

HUEBNER, R. E., SCHEIN, M. F. & BASS JR, J. B. 1993. The tuberculin skin test. *Clinical Infectious Diseases***,** 968-975.

ISTEPANIAN, R., LAXMINARAYAN, S. & PATTICHIS, C. S. 2006. *M-health*, Springer, 3-14.

KIM, Y.-T. 1997. Contrast enhancement using brightness preserving bi-histogram equalization. *IEEE transactions on Consumer Electronics,* 43**,** 1-8.

KRITZINGER, F. E., DEN BOON, S., VERVER, S., ENARSON, D. A., LOMBARD, C. J., BORGDORFF, M. W., GIE, R. P. & BEYERS, N. 2009. No decrease in annual risk of tuberculosis infection in endemic area in Cape Town, South Africa. *Tropical Medicine & International Health,* 14**,** 136-142.

LINAS, B. P., WONG, A. Y., FREEDBERG, K. A. & HORSBURGH JR, C. R. 2011. Priorities for screening and treatment of latent tuberculosis infection in the United States. *American journal of respiratory and critical care medicine,* 184**,** 590-601.

LOPEZ, A. D., MATHERS, C. D., EZZATI, M., JAMISON, D. T. & MURRAY, C. J. 2006. Global and regional burden of disease and risk factors, 2001: systematic analysis of population health data. *The Lancet,* 367**,** 1747-1757.

LOSA, G. A. 2011. *Fractals in biology and medicine*, Wiley Online Library.

MAIER, T., KULICHOVA, D., SCHOTTEN, K., ASTRID, R., RUZICKA, T., BERKING, C. & UDREA, A. 2015. Accuracy of a smartphone application using fractal image analysis of pigmented moles compared to clinical diagnosis and histological result. *Journal of the European Academy of Dermatology and Venereology,* 29**,** 663-667.

MARDANI, M. & ABTAHIAN, Z. 2015. New Advances in Diagnosis of Latent Tuberculosis Infection: A Review Article. *Archives of Pediatric Infectious Disease;* 3(1 TB):e22368.

MEYER, S. N., HOUGEN, A. & EDWARDS, P. 1951. Experimental error in the determination of tuberculin sensitivity. *Public Health Reports (1896-1970)***,** 561-569.

MOGHADDAM, R. F. & CHERIET, M. 2012. AdOtsu: An adaptive and parameterless generalization of Otsu's method for document image binarization. *Pattern Recognition,* 45**,** 2419-2431.

MORÁN-MENDOZA, O., TELLO-ZAVALA, M., RIVERA-CAMARILLO, M. & RÍOS-MEZA, Y. 2013. Comparison of different methods and times for reading the tuberculin skin test. *The International Journal of Tuberculosis and Lung Disease,* 17**,** 1273-1278.

OLIVEIRA, S. M. D. V. L., RUFFINO-NETTO, A., PANIAGO, A. M. M., OLIVEIRA, O. A. D., MARQUES, M., CUNHA, R. V. D. & ANDREOTTI, R. 2011. Tuberculin skin test:

operational research in the state of Mato Grosso do Sul, Brazil. *Jornal Brasileiro de Pneumologia,* 37**,** 646-654.
ORTBLAD, K. F., LOZANO, R. & MURRAY, C. J. 2013. The burden of HIV: insights from the Global Burden of Disease Study 2010. *Aids,* 27**,** 2003-2017.
OTSU, N. 1975. A threshold selection method from gray-level histograms. *Automatica,* 11**,** 23-27.
PHOTOMODELER. 2012. PhotoModeler. Available: http://www.photomodeler.com/ [Accessed 12 October 2016].
POUCHOT, J., GRASLAND, A., COLLET, C., COSTE, J., ESDAILE, J. M. & VINCENEUX, P. 1997. Reliability of tuberculin skin test measurement. *Annals of internal medicine,* 126**,** 210-214.
RANGAKA, M. X., WILKINSON, K. A., GLYNN, J. R., LING, D., MENZIES, D., MWANSA-KAMBAFWILE, J., FIELDING, K., WILKINSON, R. J. & PAI, M. 2012. Predictive value of interferon-γ release assays for incident active tuberculosis: a systematic review and meta-analysis. *The Lancet infectious diseases,* 12**,** 45-55.
SALAZAR-GAMARRA, R., SEELAUS, R., DA SILVA, J. V. L., DA SILVA, A. M. & DIB, L. L. 2016. Monoscopic photogrammetry to obtain 3D models by a mobile device: a method for making facial prostheses. *Journal of Otolaryngology-Head & Neck Surgery,* 45**,** 1.
SCHENK, T. 2005. Introduction to photogrammetry. *The Ohio State University, Columbus*.
SKLANSKY, J. 1978. Image segmentation and feature extraction. *IEEE Transactions on Systems, Man, and Cybernetics,* 8**,** 237-247.
SMITH, M. & KOKKAS, N. Assessing the photogrammetric potenctial of cameras in portable devices. International Archives of the Photogrammetry, In: XXII ISPRS Congress of the International Society for Photogrammetry and Remote Sensing, 2012. B5.
SUDHEER, P. 2016. Wound Assessment System for Patients with Diabetes Identification. *International Journal of Advance Research in Computer Science and Management Studies,* 84.
SUROVÝ, P., YOSHIMOTO, A. & PANAGIOTIDIS, D. 2016. Accuracy of Reconstruction of the Tree Stem Surface Using Terrestrial Close-Range Photogrammetry. *Remote Sensing,* 8**,** 123.
SWAHN, M. H., BRAUNSTEIN, S. & KASIRYE, R. 2014. Demographic and psychosocial correlates of mobile phone ownership and usage among youth living in the slums of Kampala, Uganda. *Western Journal of Emergency Medicine,* 15.
TRAJMAN, A., STEFFEN, R. & MENZIES, D. 2013. Interferon-gamma release assays versus tuberculin skin testing for the diagnosis of latent tuberculosis infection: an overview of the evidence. *Pulmonary medicine,* 2013.
UN 2014. World Urbanization Prospects: The 2014 Revision, Highlights (ST/ESA/SER. A/352). *New York, United*.
VARMA, A., VARMA, A., JAKKAMPUDI, H. M. & RAJ, A. N. J. 2016. Vision based Pus Segmentation and Area Estimation of Wound using Android Application. *Indian Journal of Science and Technology,* 1-9.
WALFORD, A. 2007. *What is photogrammetry?* [Online]. Available: http://www.photogrammetry.com/ [Accessed 15 October 2016].
WALLIS, L. A., FLEMING, J., HASSELBERG, M., LAFLAMME, L. & LUNDIN, J. 2016. A smartphone app and cloud-based consultation system for burn injury emergency care. *PloS one,* 11**,** e0147253.

WHO 2006. Global tuberculosis control: surveillance, planning, financing: WHO report 2006, World Health Organisation.
WHO 2013. *Global tuberculosis report 2013*, World Health Organization.
WHO 2015. *Global tuberculosis report 2015*, World Health Organization.
WILSON, N., COUPER, I., DE VRIES, E., REID, S., FISH, T. & MARAIS, B. 2009. A critical review redress the inequitable distribution of healthcare professionals to rural and remote areas. *Rural and Remote Health,* 1-21.
WU, C. 2011. Visualsfm: A visual structure from motion system (2011). *URL* http://www/*. cs. washington. edu/homes/ccwu/vsfm,* 14 [Accessed 12 October 2016].

11. APPENDICES

11.1 Appendix A – Scaling sticker design

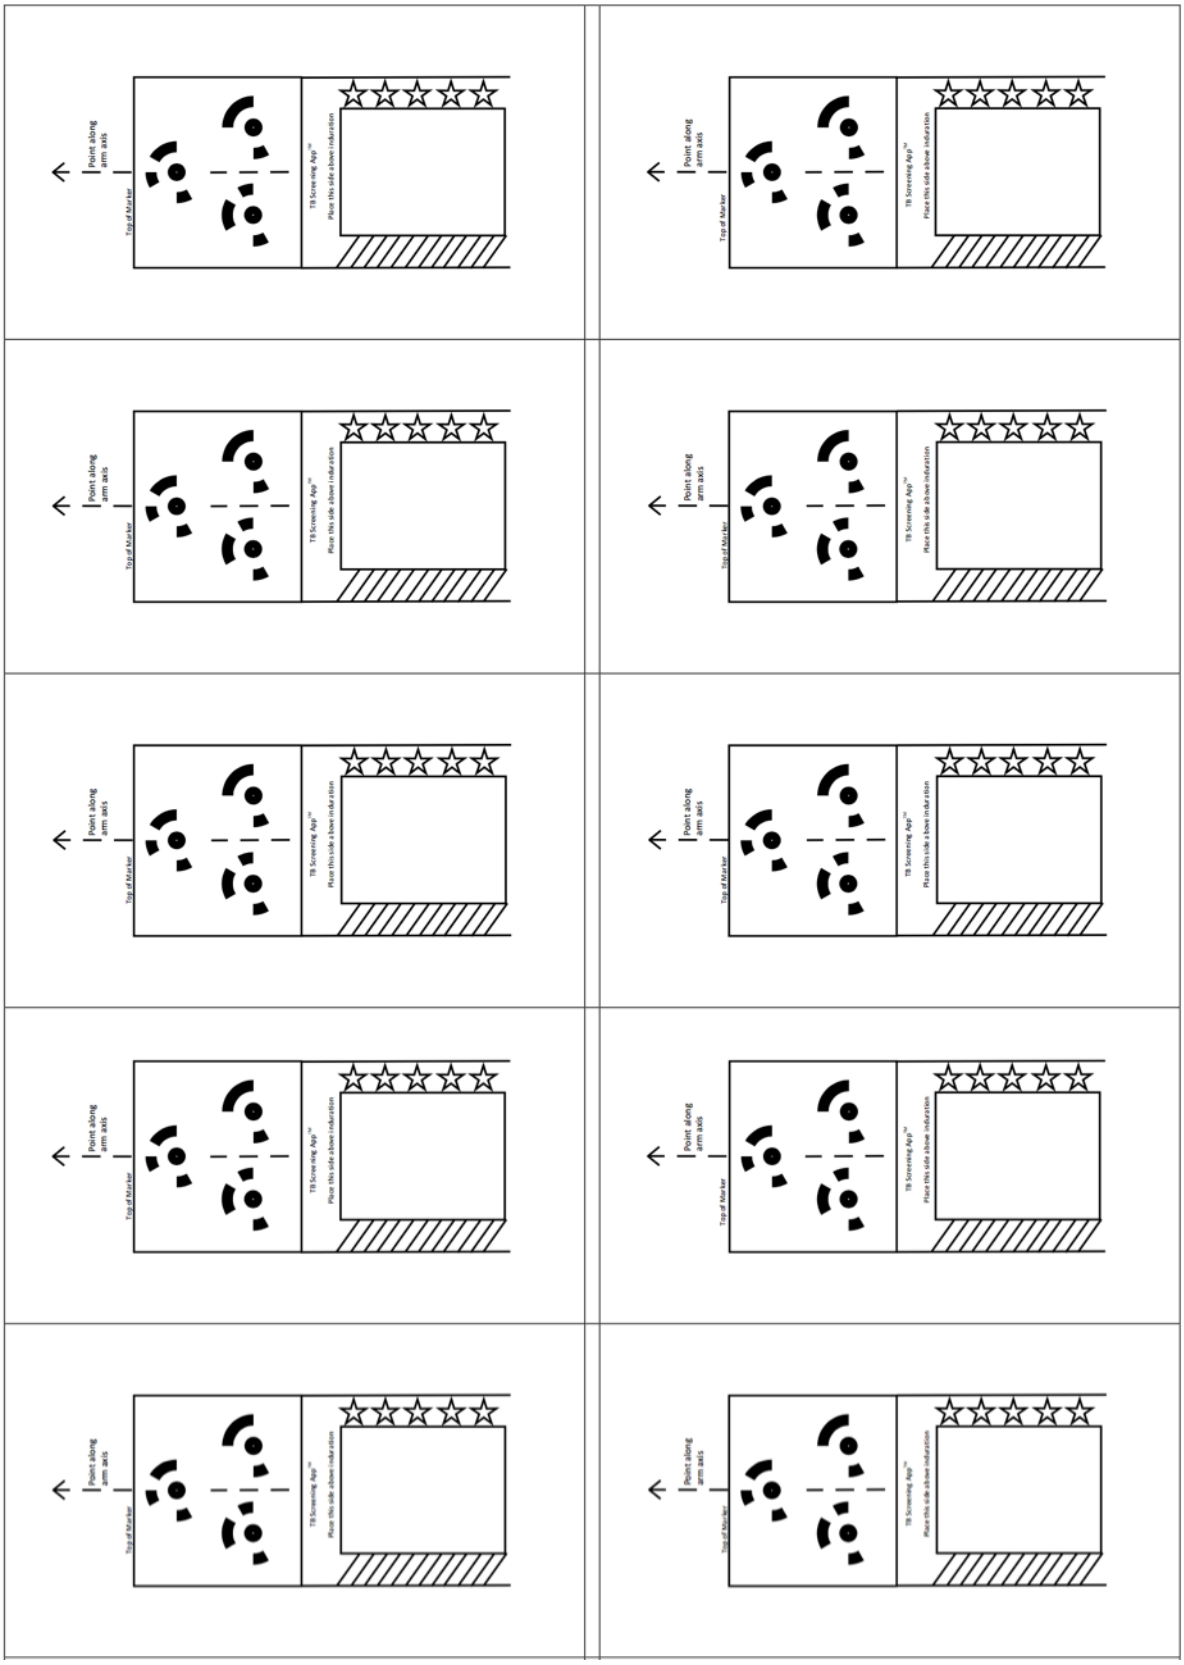

11.2 Appendix B – TB screening app evaluation form

TB Screening App Evaluation Form

Please answer the following questions regarding the usability of the TB Screening App:

		Strongly Disagree				Strongly Agree
1.	I think that I would like to use this mobile app frequently.	☐	☐	☐	☐	☐
2.	I found this mobile app unnecessarily complex.	☐	☐	☐	☐	☐
3.	I thought this mobile app was easy to use.	☐	☐	☐	☐	☐
4.	I think that I would need assistance to be able to use this mobile app.	☐	☐	☐	☐	☐
5.	I found the various functions in this mobile app were well integrated.	☐	☐	☐	☐	☐
6.	I thought there was too much inconsistency in this mobile app.	☐	☐	☐	☐	☐
7.	I would imagine that most people would learn to use this mobile app very quickly.	☐	☐	☐	☐	☐
8.	I found this mobile app very cumbersome/awkward to use.	☐	☐	☐	☐	☐
9.	I felt very confident using this mobile app.	☐	☐	☐	☐	☐
10.	I needed to learn a lot of things before I could get going with this mobile app.	☐	☐	☐	☐	☐

Do you have any suggestions on how the app could be improved? If yes, please share below:

_____
_____
_____
_____
_____
_____

Do you have any other comments regarding the application? If yes, please share below:

_____
_____
_____
_____
_____

*Questions in this questionnaire were adapted from online questionnaire found at http://satoriinteractive.com/system-usability-scale-a-quick-usability-scoring-solution/