Hydration sensor development for use on infants in underserved areas

by

Eduard Kieser

Thesis presented in partial fulfilment of the requirements for the degree of Master in Engineering (Mechatronic) at the University of Stellenbosch

Supervisor: Dr Dawie van den Heever

Co-supervisor: Prof Kiran Dellimore

December 2015

Declaration

By submitting this thesis electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Date: September 2015 .....

Abstract

Hydration sensor development for use on infants in underserved areas

E. Kieser

Department of Mechanical and Mechatronic Engineering,
University of Stellenbosch,
Private Bag X1, Matieland 7602, South Africa.

Thesis: MEng Research (Mechatronic)

December 2015

To treat dehydration in infants a good measure of dehydration level is required. Clinical hydration scales are the standard for assessing hydration upon first examination. These scales rely on the training of the physician and do not perform well when applied in underserved settings where the vast majority of infant mortalities due to dehydration occur. This study proposes three sensors for the non-invasive assessment of dehydration. All three sensors measure proven indicators of hydration (capillary refill time, skin recoil time and skin temperature difference) and the aim is to eventually allow untrained personnel to reliably measure these hydration markers with greater inter-rater reliability. A proof of concept prototype of each sensor was fabricated and the sensors were tested in two separate studies. The skin recoil time and temperature profile sensors seemed to perform fairly well when tested on infants, however, because of the small sample size of the infant study, no real conclusions could be drawn about the link between the measured markers and hydration state. A systematic method of finding a model, that combines the best predictors, was presented, however the model performed only marginally better than the best individual marker, most likely due to wanting orthogonality between the predictors. Skin recoil time, which is the marker that performed best in the clinical study, achieved a sensitivity and specificity of 0.8 and 0.84 respectively for detecting 5% dehydration, which is better than the expected performance of any single commonly accepted hydration marker. Interestingly, this combination of sensitivity and specificity was achieved at a binary threshold of 0.08 seconds. As mentioned previously, the data does not make it possible to draw solid conclusions about the performance of the device, however the

presented evidence makes a strong case for continued research into the presented methods.

Uittreksel

Hidrasie-sensor ontwikkeling vir gebruik by kleinkinders in onderbediende areas

E. Kieser

Departement Meganiese en Megatroniese Ingenieurswese,
Universiteit van Stellenbosch,
Privaatsak X1, Matieland 7602, Suid Afrika.

Tesis: MIng Navorsing (Megatronies)

Desember 2015

Dit is belangrik om dehidrasie in kleinkinders akkuraat te meet om sodoende die regte behandeling toe te pas. Kliniese hidrasie skale is die standaard vir die bepaling van hidrasie met die eerste ondersoek. Hierdie skale vertrou op die opleiding en ondervinding van die dokter en is nie baie betroubaar wanner dit gebruik word in onbediende areas waar die oorgrootte meerderheid van kindersterftes weens dehidrasie is. Hierdie navorsing stel drie sensors voor vir nie-indringende assessering van dehidrasie. Al drie sensors meet bewese indikators van dehidrasie (kapillêre hervultyd, vel-terugkrimping en veltemperatuur-verskil) en die doel is om uiteindelik onopgeleide personeel toe te laat om hierdie merkers betroubaar te meet. ‘n Bewys-van-konsep prototipe van elke van die sensors is vervaardig en getoets in twee verksillende studies. Die vel-terugkrimping en veltemperatuur-verskil sensors het heel goed gevaar toe dit op die kleinkinders getoets is, maar vanweë die klein toetsgroep kon geen ware afleidings gemaak word tussen die verwantskap van die gemete merkers en die hidrasietoestand nie. ‘n Sistematiese metode om ‘n model te vind wat die beste voorspellers kombineer is voorgestel, maar die model het egter slegs marginaal beter presteer as die beste individuele merker, heel moontlik vanweë ‘n tekort aan ortogonaliteit tussen die voorspellers. Vel-terugkrimping, die merker wat die beste presteer het in die kliniese studie, het onderskeidelik ‘n sensitiwiteit en spesifiekheid van 0.8 en 0.84 behaal vir die opsporing van 5% dehidrasie, wat beter is as die verwagte verrigting van enige van die ander algemeen aanvaarde hidrasiemerkers. Interessant genoeg is die kombinasie van sensitiwiteit en spesifiekheid behaal by ‘n binêre grens van 0.08 sekondes. Soos voorheen genoem is dit met die data nie moontlik om konkrete gevolgtrekkings

te maak rakende die gedrag van die toestel nie, maar die voorgestelde bewyse maak ‘n sterk saak vir verdere navorsing oor die voorgestelde metodes.

Dedications

Ek wil graag hierdie werk opdra aan Prof Scheffer, wie se geduld en ondersteuning dit alles moontlik gemaak het.

Nomenclature

Abbreviations

ADC	Analogue to digital converter
AUC	Area under the receiver operating characteristic curve
BMI	Body mass index
CRT	Capillary refill time
dCRT	Digitally measured capillary refill time
DEA	Data extraction application
DOF	Degree(s) of Freedom
dSRT	Digitally measured skin recoil time
ECF	Extra cellular fluid
FFM	Fat free mass
FoV	Field of view
HREC	Health research ethics committee of Tygerberg
ICF	Intra cellular fluid
KDE	Kernel density estimate
LBM	Lean body mass
LDA	Linear discriminant analysis
OCR	Optical character recognition engine
PCA	Principle component analysis
P_{osm}	Plasma osmolality
ROC	Receiver operating characteristic
ROI	Region of interest
SRT	Skin recoil time
TBW	Total body water
TP	Thermal profile
WHO	World health organization

Contents

List of Figures

List of Tables

1. Introduction

1.1 Project motivation

There are several aspects of this project that require motivation. The following section explains briefly why the project is important and goes into some of the rationale that lies behind the most important decisions that were made.

The ability to conveniently and reliably measure the hydration status of a patient is a significant unmet medical need [1].

Diarrhoea is the most prevalent childhood disease in all regions of the world. It is estimated that diarrhoea and the associated dehydration kills roughly 1.9 million children annually, which accounts for 19% of all deaths in children under the age of five [2].

Experienced physicians can generally make fairly secure assessments and ensure that the patients receive the appropriate treatment, but well trained physicians are scarce in poor and rural areas where diarrhoea and dehydration is most prevalent [3]. The under- or overestimation of hydration state in these cases can lead to the unnecessary use of re-hydration supplies and can, in the worst cases, lead to preventable infant deaths.

A system that can assist in the assessment of hydration state of infants, in a safe and non-invasive way, in a rural setting, is therefore very desirable.

1.2 Context, scope and clarification of parallel projects

Philips Research has identified the need for a hydration assessment system and has agreed to fund two Stellenbosch based projects to pursue the development of such a system. This project will focus on the use of existing technology and will aim to develop sensors that have reasonable potential to be industrialized in such a way that they can operate in a rural setting. While the basic technology may not be novel, the use of it for the assessment of hydration level is. This project will therefore not aim to produce market ready designs. Instead, some sensor concepts will be proposed and prototypes of the concept designs will be built to test the proposed concepts as thoroughly as possible. The information from the tests will then be used to comment on the potential of the concepts

and to make recommendations for future projects. As stated earlier this project runs in parallel (and at some parts overlaps) with that of Mr. Cobus Visser who is an Electric and Electronic Engineering graduate from Stellenbosch, who is also pursuing a Master's degree at the Biomedical Engineering Research Group.

1.3 Project goals

The above information can be distilled down to a concise set of project goals. The over-arching theme of which, is the reduction of infant mortality due to diarrhoea associated dehydration.

- Propose a device concept or system concept that can:
  - Assist assessors of dehydration by providing a quantitative basis for the markers that are used for infant hydration assessment.
  - Provide hydration markers which have a diagnostic value that is similar to, or better than, that of conventional markers that are measured by the physicians as quantified by predictor sensitivity and specificity for a specific level of dehydration.
- Develop a set of prototypes that:
  - Can mimic the operation of the proposed concepts to such an extent that they may be used to assess the value of the proposed devices.
  - Will be ergonomic enough to use in hospital setting so that they may be tested on dehydrated patients.
- Evaluate the potential of the proposed concepts by examining the performance of the prototypes.
- Evaluate the feasibility of refining the devices to the point where they will be useful in the field.

2. Literature Review

2.1 What is dehydration?

Very loosely stated, dehydration is a state that occurs when a person loses more fluids than they take in, to such an extent that the loss in body fluid starts to affect their normal metabolic processes. The level of dehydration is generally expressed as the percentage loss in body mass due to fluid loss.

2.2 Normal body water composition

Around 82% of the lean body mass (LBM) also called fat free mass (FFM) of term newborns is made up of water. This is known as the total body water (TBW) of the infant. This value decreases to around 70% by the age of four months, where it reaches a steady state [4]. In infants of around 4 months 25% of the LBM is in the extracellular fluid (ECF) and 45% of the LBM is found in the intracellular fluid (ICF). This relative TBW composition shifts slowly until the age of around 5 years when it too reaches its normal values. At this stage the ECF makes up around 20% of the LBM and the ICF accounts for 50% of the LBM. Also found in the ECF is the plasma which should make up about 6% of the LBM at all ages [5]. Therefore, in order to quantify the hydration level of infants by looking at body water composition, it is important to always remember to relate a specific measurement to a specific age as the normal values for these parameters change significantly in the first few months of a child's life.

2.3 The major types of dehydration

Dehydration or hypovolemia is the excessive decrease in total body water (TBW). The most common form of dehydration is called isonatremic (isotonic) dehydration. Isonatremic dehydration occurs when water and salt are lost in the same relative proportions. This type of dehydration is often associated with diarrhoea and vomiting and accounts for roughly 80% of all dehydration cases [6]. Hypernatremic (hypertonic) dehydration, occurs when the patient loses a large amount of fluid, but a relatively small amount of salts and is usually

associated with excessive sweating, fluid-restriction and osmotic diarrhoea [5, 6]. Hyponatrimic (hypotonic) dehydration is encountered when the patient loses fluid with relatively high sodium content. Hyponatrimic dehydration also occurs when fluid loss is ongoing and the only fluid intake is water [5].

Dehydration therefore has two major components. The first is the change in body fluid volume which is usually studied by examining the TBW or the change in patient weight. The second component of dehydration is the change in concentration of the remaining body fluid. The body fluid that provides the best indication of dehydration, based on osmolality and protein concentration is plasma [7, 8, 9]. Despite the important distinction between the different types of dehydration, the conditions are most commonly still collectively described by percentage weight loss due to water loss.

2.4 State of the art

When discussing the way that dehydration is assessed in practice we need to distinguish between the clinical methods that are used in most cases by physicians to assess hydration state and the more invasive methods that are used in research and only in very special cases for diagnosis.

2.4.1 Laboratory hydration assessment techniques

Isotope dilution is perhaps the most invasive hydration assessment technique that will be discussed here, however it is considered by experts to be the gold standard for TBW assessment[8]. The rationale behind it is fairly simple: One starts by adding a known quantity of some foreign compound to the blood stream (in most cases heavy water, D_2O). After enough time has passed for the foreign compound to homogeneously mix with the body fluids, a small sample of blood is drawn. If one then assumes that the same ratio between fluid and compound exist in both the patient and the sample, then the patient fluid volume can be calculated using:

$$body\ fluid = \ total\ amount\ of\ compound\ added \times \frac{fluid\ in\ sample}{compound\ in\ sample}, \quad (2.1)$$

where body fluid volume is the only unknown.

Bio Impedance Analysis is another method of determining TBW that has gained popularity in recent years. The method involves measuring the impedance between two places on the patient and making some deductions based on those assessments.

Freeze point tests are an established method of determining plasma osmolality (P_{osm}). Osmolality, protein concentration and conductivity analysis can

be done on many body fluids, which include blood plasma, saliva and urine. According to Cheuvront *et al.* [10], P_{osm} is the only hydration marker with the potential to be used for static dehydration assessment (one measurement). P_{osm} is a measure of the body's electrolyte water balance. P_{osm} is determined by taking a blood sample and using the freezing point to determine the amount of electrolytes in solution [8]. The higher the P_{osm} the lower the freezing temperature of the plasma component of the blood sample. Assessment of hydration by considering P_{osm} is also complicated by the fact that the normal values for these parameters vary between patients. Therefore in order to use P_{osm} to determine hydration status one would need to know the baseline value of P_{osm} for that specific patient. Fortunately this variation is to some degree systematic. This systematic component of the variation in baseline values can be related to patient age, gender and regional factors such as diet and climate [8]. ECF osmolality and therefore P_{osm} is regulated in mammals, with an average set point of roughly 287 mosm/kg in humans [10, 11]. Despite this regulation it is normal for the P_{osm} to vary around the set point with perturbations of around 1% to 3% due to intermittent changes in water loss and fluid intake as well as Na^+ intake and excretion [11].

The body regulates P_{osm} quite closely, the reason for this is that the relative osmolality between the ECF and ICF controls osmosis between the ECF and ICF. For example if the P_{osm} of the ECF were to suddenly increase, water would move from the ICF to the ECF decreasing the volume of the affected cells. Since the brain is housed in the rigid cranium, the sudden change in tissue volume has the potential to cause severe nerve damage [11].

2.4.2 Clinical hydration assessment methods

In the 2004 paper "Is this child dehydrated" Steiner et. al. [12] reviewed the results from 1561 potential articles that measure the performance of clinical markers for detecting dehydration in infants and young children. From the articles they identified 13 that contained original data, reported accepted diagnostic markers and were not limited somehow by the study design. From this review Steiner et. al. concluded: (1) The initial assessment of dehydration in young children and infants should be based on capillary refill time (CRT), skin recoil time (SRT) and respiratory patterns. (2) Combinations of signs perform better than individual signs. (3) The relative inaccuracy of currently available methods limit the clinician's ability to estimate the exact degree of dehydration.

In 1997 Gorelick [13] investigated the validity and reliability of ten popular clinical signs that are used for the hydration assessment of children. The study focussed on children between the ages of one month and 5 years. When comparing SRT, CRT, general appearance, absence of tears, abnormal respiration, dry mucous membranes, sunken eyes, abnormal radial pulse, tachycardia and decreased urine output he found that prolonged skin recoil time and capillary

refill time had the best positive predictive value for assessing the presence of 10% dehydration. He then selected a subset of four predictors (CRT, dry mucous membranes, absent tears and general appearance) which where independently associated to hydration level and conducted a logistic regression. The p-values for the components of the logistic regression were 0.0002, 0.0079, 0.0077 and 0.0427 for CRT, dry mucous membranes, absent tears and general appearance respectively. This indicated that, of the subset of predictors, CRT is by far the best correlated to hydration level. Gorelick then presented two scales for the assessment of hydration state, one based on a combination of all ten predictors and one based on a combination of the subset of four predictors. These scales subsequently became known as the Gorelick 10 and 4 point scales.

In 2004 Friedman [1] published a somewhat similar study, this time with the explicit objective of developing a dehydration scale that is based on a combination of commonly used clinical hydration markers. For the development of his scale Friedman considered: Urine output, general appearance, CRT, SRT, fontanelle appearance, eye appearance, breathing, heart rate, systolic blood pressure and urine specific gravity. All the predictors were rated on a 3-point like scale (0,1,2) to determine the marker intensity. He then proposed an 8 point scale based on the assessment of general appearance, eyes , mucous membranes and tears.

A third dehydration scale that is also commonly used for the assessment of infant dehydration is known as the WHO (World Health Orginization) dehydration scale. The WHO scale relies on the assessment of general appearance, eyes, thirst and SRT.

Finally a distinction should be made between the signs that can be measured instantly like CRT and SRT, and the ones that can only be measured over time like urine output, thirst and general appearance. Experts prefer using the latter group for accurate dehydration assessments, but often triage assessments are required on the spot which means that the latter group of predictors cannot be used. This means that while whole patient measurements are currently considered to be more accurate it would be more valuable to the clinicians to try to develop an assessment that is based on the instantly measurable signs.

3. Sensor Design

3.1 Selecting markers

The literature review discussed many innovative methods assessing hydration, all of which are continuously scrutinised and developed by capable researchers, to improve the accuracy of hydration assessment in athletes, infants and the elderly.

The sobering conclusion that Pringle and Falszewska [14, 15] made is that the currently accepted methods of clinical infant hydration assessment, which work well when executed by physicians in high-income countries, break down when used by their less specialised counterparts in underserved areas. This is especially alarming when one considers that dehydration related mortality is far more prevalent in underserved areas [3].

It is therefore in-line with the project objectives to discuss why the current systems do not work as well in underserved areas and to address this with the resources at our disposal.

It stands to reason that there are two possible explanations for why the dehydration scales would work in the developed world and not in the developing world. Either the way in which dehydration manifests is different, or the scales are used less efficiently in the developing world. Pringle holds that the latter is more likely the case and highlights the need for better hydration assessment tools, especially for infants in developing countries [14].

The three most common hydration scales all rely on manually obtained data. This means that the integrity of the data depends on the training of the physician that collects it. If one could find a way of removing some of the error involved in collecting the data, that would work in a third world setting, it would surely be a significant step in the right direction.

A prolonged capillary refill time is a sign of reduced peripheral perfusion (circulation). A large core-peripheral temperature gap has also been associated with reduced peripheral perfusion [16] and is in some cases used as a clinical marker for assessing dehydration. Intuitively it makes sense that lower peripheral perfusion would lead to a greater core peripheral temperature gap since the blood flow that is responsible for transporting the heat from the central areas of high metabolic activity to the body's extremities is reduced. One could take this one step further: We know that dehydration is linked to reduced peripheral

perfusion, we also know that reduced perfusion leads to a less even temperature distribution, it would therefore be reasonable to expect that one would see larger temperature gradients between areas of high and low metabolic activity. This is shown graphically in Figure 3.1

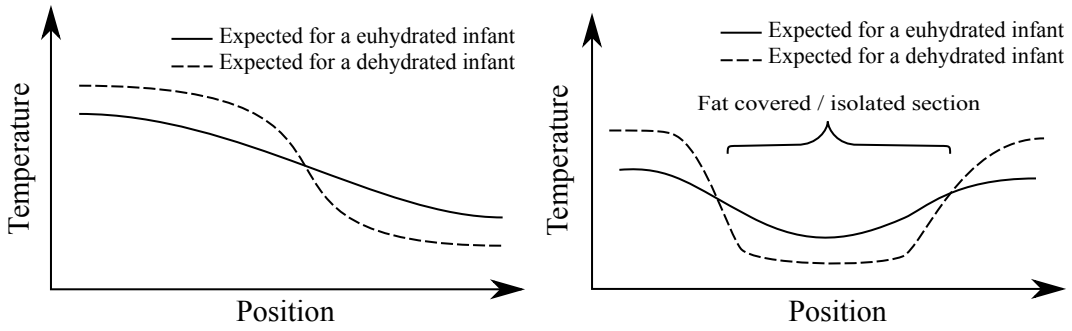

Figure 3.1: Diagramatic representation of the expected temperature profile for dehydrated and euhydrated infants. Between the core and extremities (left), over fat covered regions (right)

The sensors that will be developed in this project were chosen such that, the markers are based on accepted clinical markers that can instantaneously be measured and that have a reasonable chance of being measured more accurately by an electronic sensor. We will therefore focus the remainder of our attention on developing sensors to improve the assessment of capillary refill time, skin recoil time and skin temperature profile.

3.2 Study overview

Before examining the detailed sensor designs, it is useful to first consider the context in which the sensors will be used. The sensors are to be proof of concept devices that will be tested in a study at Tygerberg Hospital. Once admitted to the study the patients are to be examined at intervals of 4 to 6 hours. During these interval examinations (which will be referred to as snapshots from here on out) a variety of patient data is collected. Each patient snapshot is to include all the clinical observations made by the attending physician, all the observations made by the nursing staff and one measurement from each of the sensors that will be developed for this project.

For the data to be useful it is important to know which measurement corresponds to which snapshot. It is therefore important that there is a secure way of labelling the data when taking the measurements. Each patient will be assigned a set of forms. The observational data will be recorded on these forms. Each form will also have a QR code printed on it, with the patient number and snapshot number encoded therein. This QR code will then be read by the devices and used to tag the data.

3.3 Capillary refill time sensor

3.3.1 CRT sensor concept

When physicians check capillary refill time, they use their index fingers to apply a small pressure to the region of interest until blanching occurs. They then quickly remove the pressure and visually determine when the blanched area is flushed again [17, 18, 19].
The variables in this process that are likely to affect the result are:

- Magnitude of the blanching pressure
- The time that the blanching pressure is applied
- The area that the pressure is applied to
- The physician's colour perception and reaction time

The digital CRT assessment device is designed to work in a similar manner. However, the device is designed in such a way that it minimizes each of the aforementioned error inducing effects.

It is also important that the sensor be easy to operate and that the test data is securely linked to the test case.

Various designs and layouts were considered. The final device concept diagram is shown in Figure 3.2.

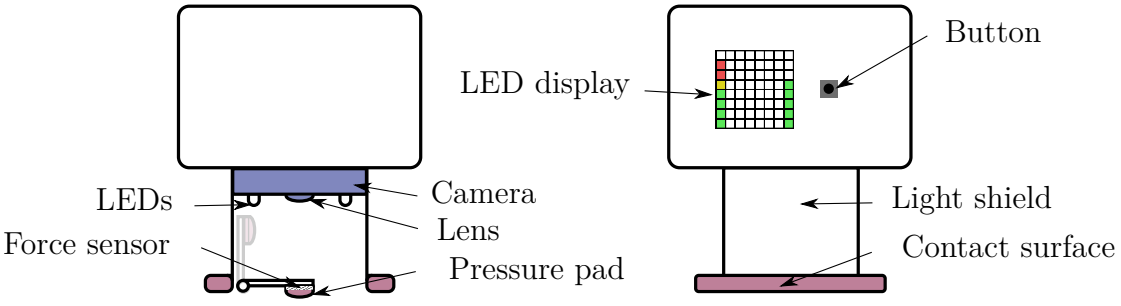

Figure 3.2: CRT device concept.

A section view of the CRT device is shown on the left of Figure 3.2. The force sensor and silicone pressure pad (which is used to apply the blanching pressure at the contact area) are both mounted on a spring loaded arm. During operation the pressure pad pushes against the contact site to cause blanching. Once the pressure has been applied for long enough the operator pulls a lever and the arm swings out of the way and allows the camera to record the capillary refill.

On the right of Figure 3.2 the front view of the CRT sensor with the control button and LED display is shown. The bar on the left of the LED display shows the measured force, which allows the operator to accurately adjust the

applied force to keep it within a pre-determined range. The bar on the right is used as a count-down timer which indicates to the operator when it is time to pull the lever that allows the arm to swing out of the way.

3.3.2 CRT sensor design

As can be see from Figure 3.2 the main moving part on the CRT sensor is the arm that holds the force sensor and pressure pad. During operation the contact surface should be resting on the patient to provide support and stability without exerting a lot of force. To achieve this the contact surface is moulded from silicone. The profile of the contact surface is designed so that it would deform to the patient's chest. It was later determined that some of the infants in the population of interest are fairly fat. This meant that the pressure pad would, in certain cases, have to protrude further past the contact surface that surrounds it. This was achieved by combining a protruding arm with a large flexible contact surface. Figure 3.3 depicts the two extreme use cases.

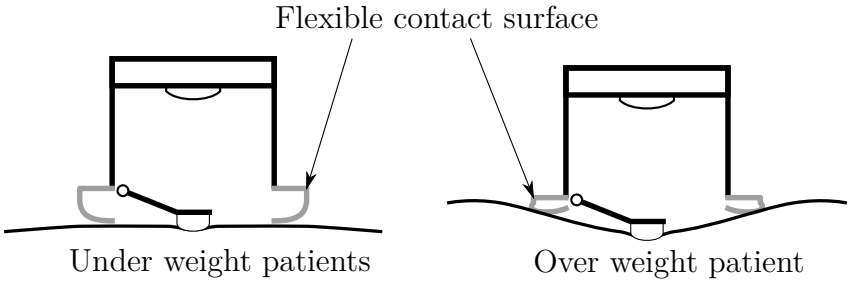

Figure 3.3: CRT sensor operation on over- and under weight patient

3.3.3 CRT sensor optics

The camera that is used in this project can only produce a clear image beyond one meter. It is therefore necessary to shorten the focal length of the camera to allow for a compact design. The simplest way of choosing the correct lens for a specific focal length is by treating the camera sensor as a far-sighted person's eye. One would then want to project a virtual image of the object that lies further away so that the camera may focus on the projected image as shown in Figure 3.4.

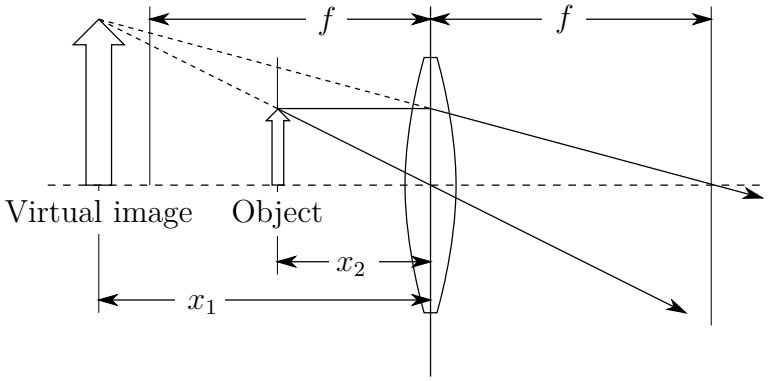

Figure 3.4: Basic thin lens parameters

Both the required distance of the Virtual image x_1 and the object distance x_2 are known. The thin lens equation as shown below can then be used to calculate the required focal length f as follows: $\frac{1}{f} = \frac{1}{x_1} + \frac{1}{x_2}$. Since the object will only be a few centimetres from the lens and the projected image should be at least one meter from the lens one can assume that $\frac{1}{x_1} \ll \frac{1}{x_2}$. The equation then simplifies to $f \approx x_2$. Based on these calculations a plano-convex lens with a 35 mm focal length was selected.

3.3.4 CRT sensor electronic hardware

To build the CRT device proposed in Section 3.3.1, a camera, a light shource, a force sensor, a display, a power source and an embedded controller that ties everything together are required. Many approaches were considered. Finally a raspberry-pi based solution was selected.

Some developers argue that the raspberry-pi is not suited for embedded applications, because of the power and computational overhead associated with running a Linux system. These are fair concerns, however, a raspberry-pi based approach offers a good balance between capable hardware, a streamlined development process and powerful low level control over the operating system and hardware settings.

The controller

The micro controller that is used in the CRT sensor is a raspberry-pi model B. It is roughly the size of a credit card and runs on a derivative of Linux. It has, a 700 MHz low power ARM applications processor, 512MB of SDRAM, two USB 2.0 ports, a 26-pin general purpose input-output header, and an Ethernet jack. While the Ethernet port does nothing to improve the functionality of the final device, it saves hours on development time. This is the case, because the Ethernet port allows the development board to connect to a local network,

which allows the developer to connect to the development board as a remote system via a capable IDE. Put simply this means that the developer can run his preferred development environment (like Eclipse) on his workstation, connect to the development board and open and edit files remotely on the micro controller. The developer can then also execute commands from his workstation that run directly on the raspberry-pi hardware and view the output on said workstation. This means that there is no 'flashing' or 'loading' of scripts, all the changes to the scripts are saved directly on the micro controller. Eclipse does however also store a local copy of all remote files, so in the case that the controller memory becomes corrupt, all is not lost.

The camera

The camera module of the raspberry pi is built around a 5 mega-pixel OmniVision OV5647 sensor. The exact sensor resolution is 2592 x 1944 pixels, it has a 53 degree horizontal and 41 degree vertical field of view. The camera module has a fixed focus lens that delivers a clear picture beyond a distance of one meter. The module connects to the raspberry pi via special serial interface.

When recording at its full sensor resolution the camera's frame rate is only 15 fps, however the camera can be set to use sub sampling and binning to achieve higher frame rates with smaller frames. The table below depicts some of the available configurations

Table 3.1: Available configurations for the raspberry pi camera. (FoV = field of view, fps = frames per second)

Resolution	Aspect Ratio	Max Framerate	FoV
2592x1944	4:3	15fps	Full
1920x1080	16:9	30fps	Partial
1296x972	4:3	42fps	Full
640x480	4:3	90fps	Full

One of the favourable properties of the OmniVision sensor, for this project, is that the exposure control, white balance and black level calibration is all handled by the raspberry-pi's GPU, which one can access via the Linux terminal. It is important that one can pre-set these parameters and disable the automatic adjustment thereof because the scene composition changes dramatically when the arm on which the pressure pad is mounted swings out of the way. The change in scene would otherwise trigger the automatic adjustment of the colour balance parameters and compromise the analytical value of the recorded data.

The display

The display in the CRT sensor is a 1.2" bi-color 8x8 LED matrix display sold by adafruit industries (NY, USA). The display has one red and one green LED per pixel for a total of 128 LEDs. When the red and green LEDs are turned on at the same lime the pixel appears orange. These LEDs are controlled via 24 pins. To simplify the use of these displays adafruit industries has developed a shield that can be used to control these pins with I^2C protocol. This allows one to control the display by writing to registers on the shield using only two digital pins. The shield uses constant current drivers to provide good control over the pixel brightness and supports dimming.

The force sensor

The force sensor that is used in the CRT sensor is a Honeywell FSS series low profile force sensor. It makes use of a piezoresistive silicon sensing element and a Wheatstone bridge for signal conditioning. The sensor is rated to handle forces of up to 15N. The sensor is roughly 9 mm long, 6 mm wide and 3 mm high (excluding the terminals).

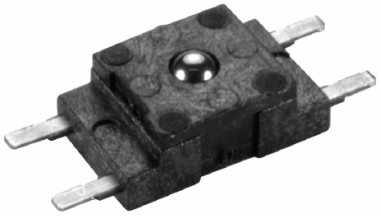

Figure 3.5: Honeywell FSS low-profile force sensor

Additional electronic hardware

The force sensor produces an analogue output, unfortunately the Broadcom chip at the core of the raspberry pi does not have any analogue inputs. This necessitates the use of an additional analogue-to-digital converter (ADC). A Texas Instruments low-power 16 bit analogue to digital converter was used to communicate between the force sensor and controller. The ADC also has a 16x gain amplifier which allows small signals like the one produced by the force sensor to be sampled with good resolution. The ADC communicates with the controller via I^2C protocol.

The CRT sensor's scene is lit by four white LEDs that are driven by a small transistor that is controlled by one of the logic outputs on the controller board. In order to keep the device as small as possible a small custom PCB was designed to house all the lighting electronics. The custom PCB also provided a convenient way to connect the ADC to the controller, which allowed for

a cleaner design. More information on the interface board can be found in Appendix B.

Power is provided by a removable power source that houses a 7.4 V 2000 mAh battery and a 5 V linear regulator.

Housing and mechanical considerations

One of the major challenges of the project was to keep the device as small as possible, without losing any of the required functionality. The final product also had to be mechanically safe (no burs or sharp edges) and sterilizable. As can be seen in Figure 3.6, the whole sensor only has three moving parts, the spring-loaded swing arm on which the force sensor is mounted, the small bracket on which the pressure pad is mounted and the latch that keeps the swing arm in place during skin blanching.

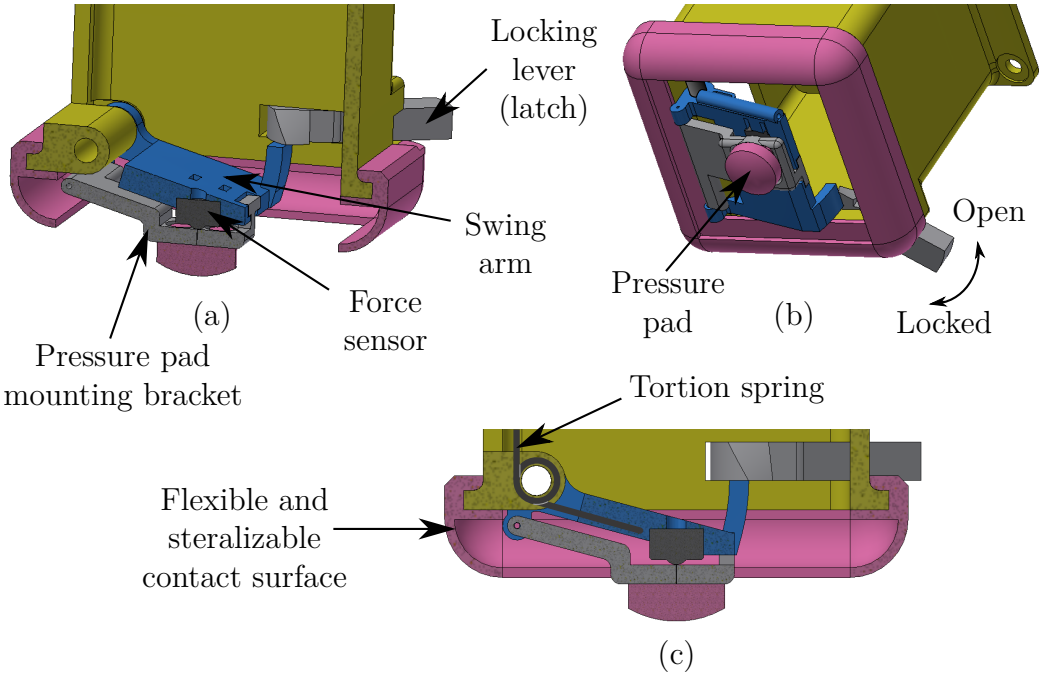

Figure 3.6: CRT sensor swing arm and locking mechanism

Despite the relatively simple mechanical operation, the design still requires complex geometries in order to meet the space constraints. The housing material also needs to be light and non conductive. The above constraints, combined with the fact that only one of each part would be produced, made 3D printing the obvious choice for the housing manufacture. All the rigid parts were therefore printed from an acrylic based material on the department's Objet 3D Printer made by Stratasys, Ltd.(Rehovot, ISR). The parts of the devices that would come into contact with the patients need to be, non reactive, bio-compatible and flexible. Therefore, all the parts that came into contact

with the patients were moulded from Body Double Standard silicon rubber from AMT-composites (Cape Town, ZAF). The final product is shown in Figure 3.7 below. Drawings of the CRT sensor and its sub components can be found in Appendix A.

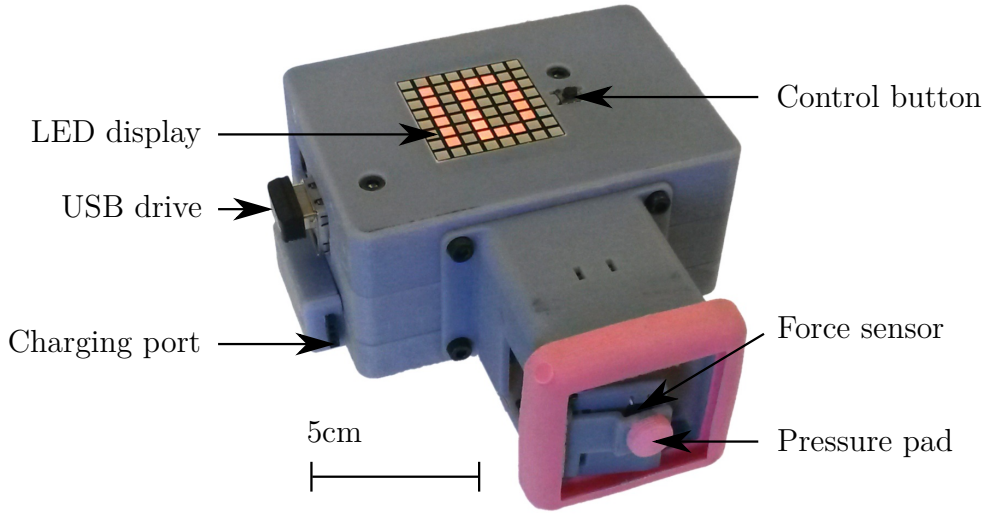

Figure 3.7: CRT sensor

3.3.5 CRT sensor firmware

Platform overview

For the CRT sensor a combination of Python and Bash script was used. Python is a concise, interpreted programming language with an extensive standard library. It supports structured and object oriented programming and is available under a GPL-compatible licence which means that is it free to use, modify and distribute.

CRT sensor work-flow

The CRT sensor control software is built as a state machine. The sensor operation can be broken down to three distinctive phases as shown below in Figure 3.8(a).

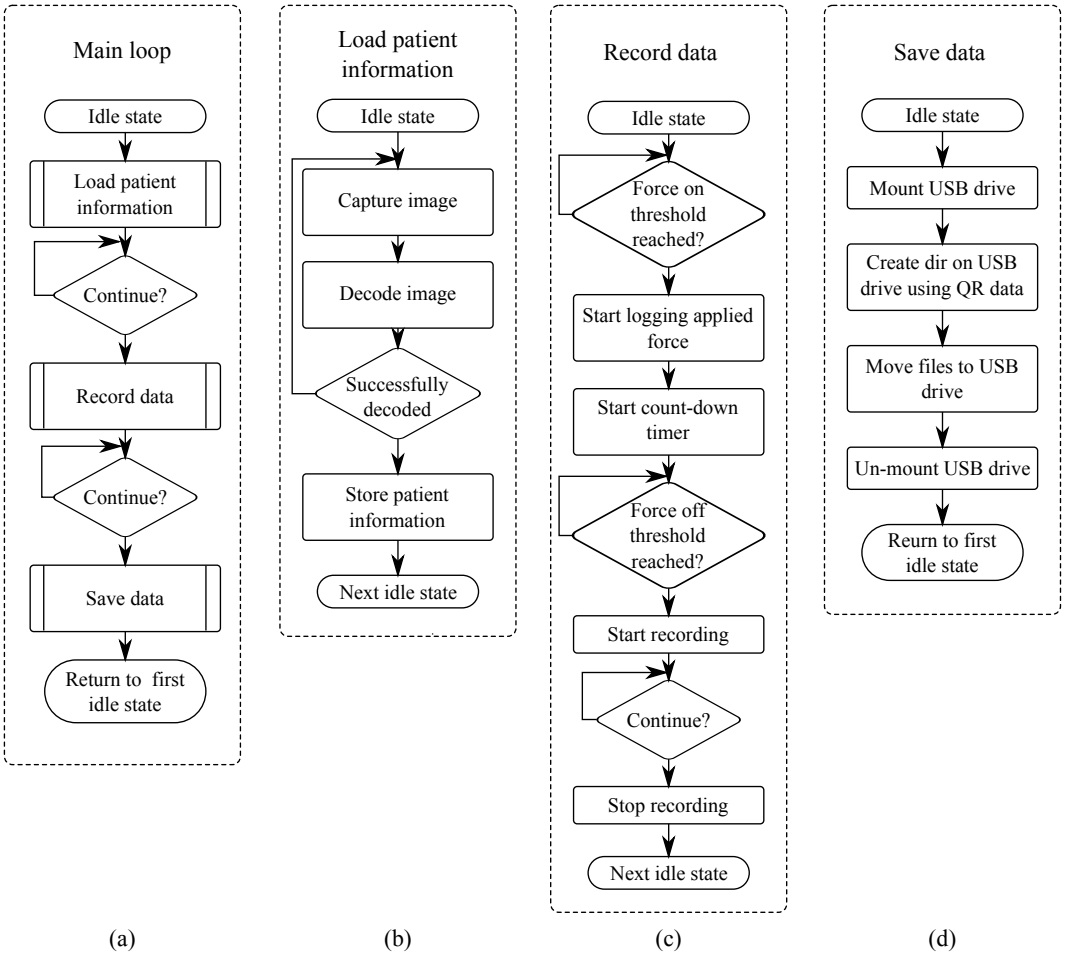

Figure 3.8: States of the CRT sensor firmware

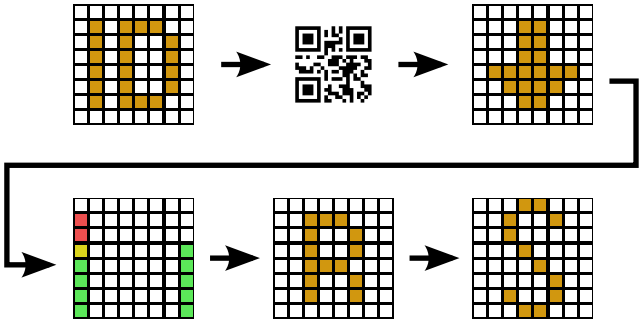

Figure 3.9: CRT sensor display during different phases of operation

The first sub-process in Figure 3.8 a, which is expanded in Figure 3.8 b represents the phase of operation where the data tags are loaded from the patient form and stored on the device. During the initial idle state the display shows the letters "ID" in orange and waits for a button push. When the process

starts the camera takes a picture, assumes that it has a picture of a QR code and tries to decode it. If the decoding fails it will repeat this process until it successfully decodes a frame. Once the information is decoded the letters on the display turn green for two seconds and the device moves to the next idle state.

Trial and error has shown that a 200x200 pixel photo of the QR image is both small enough to quickly be decoded by the on-board processor and large enough to capture all the information encoded in the QR code.

At this point one might be concerned that, if the QR code is damaged or smudged that it might be incorrectly read. Fortunately this is not the case. QR codes are based on Reed-Solomon error-correcting codes which have check symbols built into the data. Therefore, in the event that the code is damaged beyond recognition the reader will return an error instead of the wrong information.

At the start of the recording phase the display on the CRT sensor shows an orange arrow that points downward. This is to indicate that the operator can now lock the swing arm into place. Once the arm is locked into place the operator should press the button again, to tell the device that the test is about to start and that is should start monitoring the applied force. At this point the arrow will turn green. The operator can now use the device to apply the blanching pressure, by positioning the device in such a way that the pressure pad is in contact with the blanching site and gently pressing down on the device. Once the predetermined force threshold is reached the device will start logging the applied force and start the countdown timer. The device will also display the applied force in the form of a bar on the left hand side of the display. This serves as feedback to the operator who should try and keep the bar in the middle of its range, (the orange pixel) as shown in the 4th frame in Figure 3.9. Once the timer (bar on the right) has run out the operator should pull the lever that holds the arm in place to allow it to swing back. The controller will then register that the force has been removed, which will trigger the camera and illumination system. The camera will then record the scene for 10 seconds, while a yellow 'R' is displayed, and switch off. The device then moves to the pre-saving idle state and displays a yellow 'S'. If the user is satisfied he can then press the button again and the data will be moved to the USB drive. The 'S' will then turn green for two seconds, after which the device will reset itself. If the user is not satisfied with the recording he can hold the control button in for longer than two seconds to restart the procedure without saving.

3.4 Skin recoil sensor

3.4.1 Skin recoil sensor concept

When a physician assesses skin turgor (as determined by recoil time) he will pinch and pull a small piece of skin, let it loose and determine how long it takes for the skin to return to it's normal position. The measurement is strongly influenced by the physician's judgement of time. Recoil time in healthy and moderately dehydrated children is generally less than two seconds [13]. Such short times are difficult to judge accurately, which has lead many hospitals to adopt more subjective recoil time scales. In one example of such a scale the classifications are: 'instant', 'slow' and 'very slow' [13]. In order to accurately assess the skin recoil time of the patients in the planned study the skin recoil of the patients would be recorded by a camera. The recoil time would then be extracted from the footage at a later stage.

Skin recoil should also be affected by how the physician manipulates the patient's skin. If the physician pulls the skin further one might expect the recoil time to be longer. However, this error might be mediated by the fact that the reaction forces on the portion of skin will be larger when pulled further. A detailed mechanical simulation of human skin is far beyond the scope of this project. Also developing a mechanical device to consistently actuate the patient's skin in the required manner will be impractical, since the device will have to adapt to various skin types.

Development therefore carried on under the assumption that the error induced by various pinching techniques would be fairly small. This assumption was verified upon completion of the device and it was found that using techniques ranging form one extreme method to the other lead to a measurement error of around 30%, which for a healthy patient with a recoil time of less then 0.5 seconds translates to a measurement error of less than 0.15 seconds.

The proposed sensor will therefore use a camera to assess recoil time, while relying on the physician/researcher to actuate the patient's skin as shown in Figure 3.10.

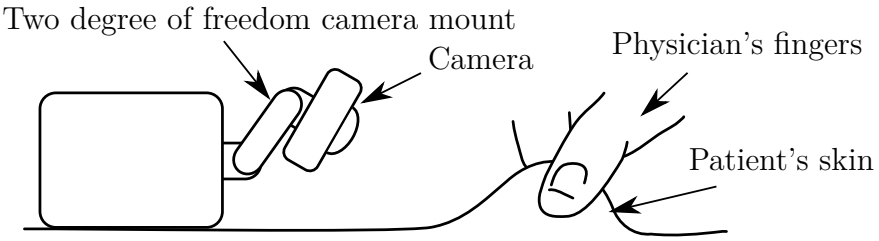

Figure 3.10: SRT device concept.

3.4.2 SRT sensor design

In order to get good data it is important that the composition of the scene is consistent. It is also important that the scene is the correct distance from the lens to ensure that the picture is in focus. Early trials with the device have shown that this is difficult to judge, without some aid. It was also found that the lighting in the wards where the devices would be tested is fairly poor. In order to address these two needs the lighting system shown in Figure 3.11 was developed. It is designed in such a way that the light from the two LEDs is funnelled into two cones that overlap at the focal length of the camera. The radii of the two cones of light at the camera's focal length is roughly the same as the camera's field of view at that length. This allows for a well lit scene while ensuring that the scene is in the camera's field of view and in focus. The three images in Figure 3.12 show how the SRT sensor's lighting system is used to frame the shot.

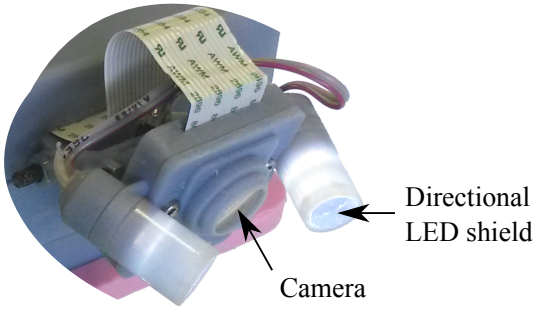

Figure 3.11: SRT sensor lighting and framing system

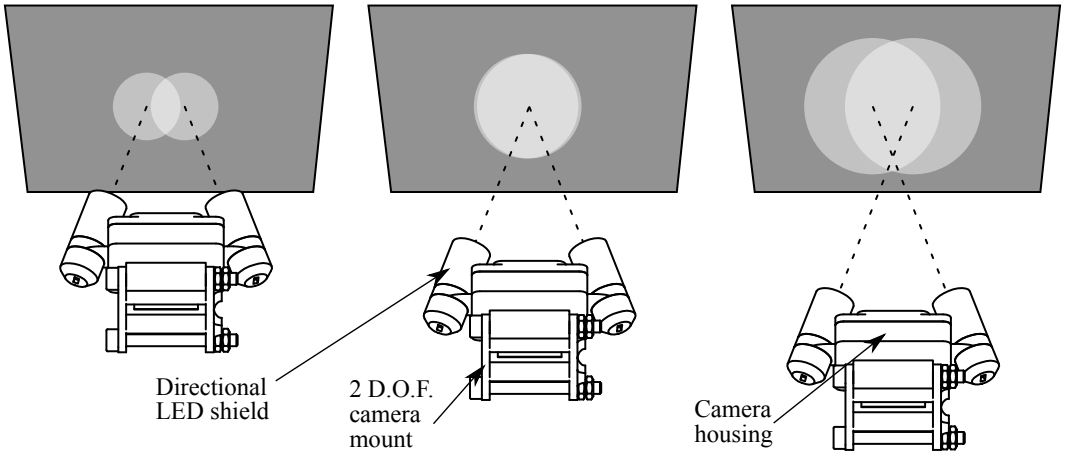

(a) Too close to scene (b) Correct distance from scene (c) Too far from scene

Figure 3.12: SRT lighting and framing system operation (D.O.F. = degree of freedom)

3.4.3 SRT sensor electronic hardware

The SRT sensor uses the same controller, camera, display and power source as the CRT sensor. In order to get a high resolution estimate of skin recoil the camera is set to record 640x480 pixel frames at 90 fps.

3.4.4 SRT sensor firmware

The control software of the skin turgor sensor is largely the same as that of the CRT sensor. In fact the only large deviation is during the data recording phase, where instead of force sensor data and timers, the control button is used to start and stop recording. Figure 3.13 briefly summarizes the SRT sensor work flow.

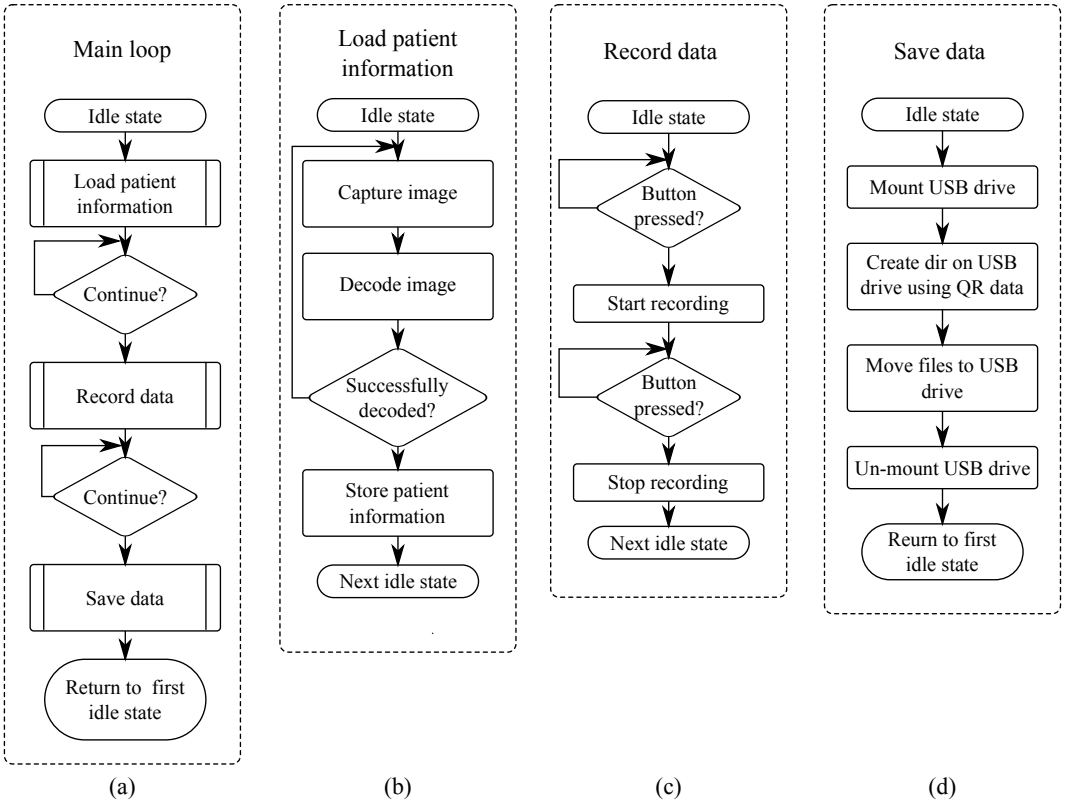

Figure 3.13: States of the SRT sensor firmware

The SRT sensor records a video of three sequential clinical skin recoil time operations. Figure 4.14 shows four sequential frames from a SRT video assessment of one recoil event. The times are then manually extracted from the video's as discussed in Chapter 4.2

3.4.5 SRT sensor housing and mechanical considerations

From a mechanical point of view the SRT sensor is very simple. The only moving parts are the two degree of freedom camera mount that can be seen in Figures 3.10 and 3.14. The joints are simple friction joints, the stiffness of which can be adjusted by tightening or loosening the connecting bolt shown in Figure 3.10. Drawings of the SRT sensor and its sub components can be found in Appendix A.

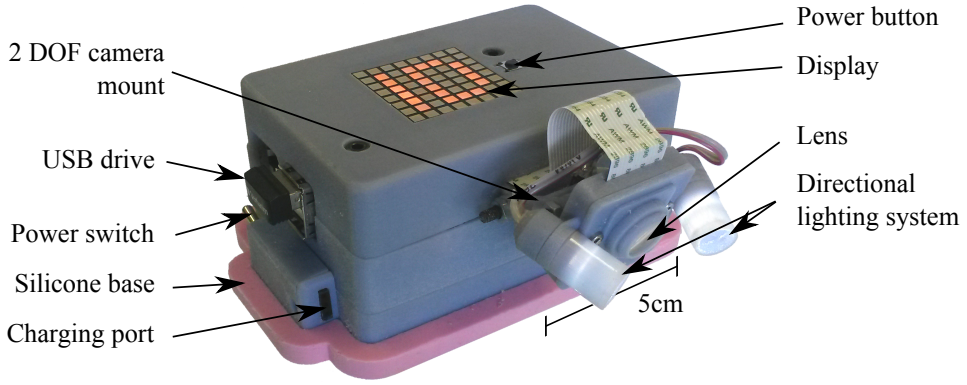

Figure 3.14: SRT sensor

3.5 Thermal profile sensor

The thermal profile sensor is intended to determine the temperature profile of a patient along a specific path, in order to validate the somewhat novel hypothesis that temperature gradient (with respect to distance) can be related to hydration level as discussed in Section 3.1. Two examples of what such a sensor might look like are shown in Figure 3.15

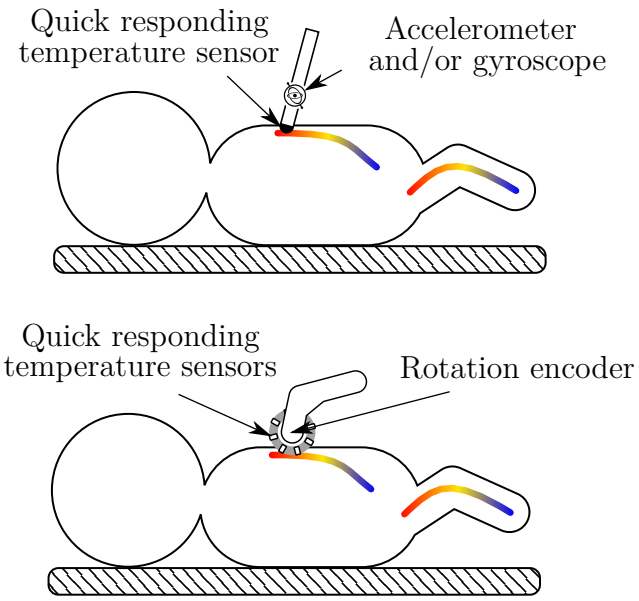

Figure 3.15: Illustrations of potential concepts for the temperature profile sensor

Since the data only needs to be collected in a laboratory/trial setting it will be far cheaper and easier to use a thermal camera to record the data and to then extract the required temperature profiles from the thermal images. This has the added benefit that it captures the temperature across all possible paths, which means that all paths can be examined during the data analysis phase.

The faculty's FLIR E60, shown below, was therefore used to collect data during the sensor concept testing at Coetzenburg and Tygerberg. The FLIR E60 has a 320x240 pixel sensor, a thermal range of $-20\,°C$ to $650\,°C$ and a thermal sensitivity of $<0.05\,°C$

Figure 3.16: An image of the FLIR E60 thermal camera that was used for the thermal profile sensor proof of concept testing

A description of the software that extracts the thermal paths from the images, in order to mimic the operation of the devices in Figure 3.15 can be found in Section 4.3

4. Marker Extraction Software

All the data that is collected by the sensors used in this project is either in the form of images or video. The next step is to take the images and videos and extract the markers that we need so that we may study how they relate to hydration level. This chapter documents the development of the application that is used to extract the desired biological markers from the study data. Figure 4.1 shows a simplified depiction of the flow of data within the data extraction application (DEA).

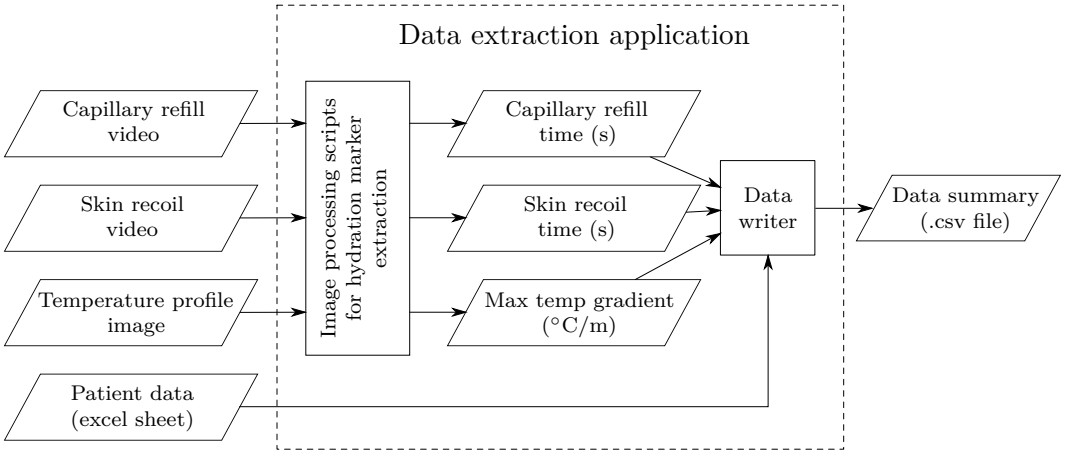

Figure 4.1: Data extraction application

The DEA is a collection of Python programs, that are tied together with a QT based user interface. The DEA is not designed to implement or test code that might be used in a final product. The DEA is developed exclusively to help the investigator to extract the required hydration markers in order to determine the clinical value of the sensor concepts. It is therefore optimized to be used with large batches of data and favours data validity over automation and computational efficiency.

4.1 Capillary refill time extraction

Before developing an algorithm to extract CRT information from the data it is instructive to first take a close look at the raw data. When doing so one will notice that there are two components to the capillary refill process. Firstly, the size of the blanching site decreases, at the same time the average colour in the remaining blanching area approaches its normal colour value.

This is best illustrated by looking at the following conditioned frames (Figure 4.2) that were adapted from data from the Coetzenburg trial, which will be discussed briefly in Section 5.1.

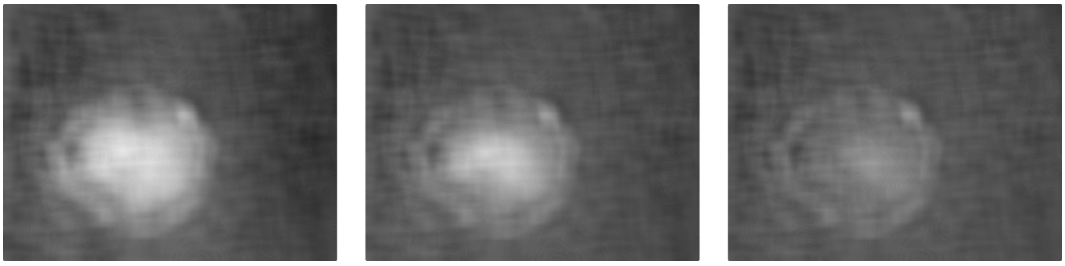

Figure 4.2: Blanching illustration, which depicts the size reduction of the blancing cite aswell as the decrease in average colour in the remaining region.

There are therefore two metrics that can be used to describe how capillary refill happens. This might seem like an overly complicated way of looking at capillary refill if the only parameter of interest is capillary refill time. However, it is important to discuss these nuances in order to develop a quantitative and objective way of determining when the capillaries have fully refilled. This will also allow for a better understanding of the limitations of the final algorithm.

The two most obvious methods that follow directly from the discussion above involve segmentation of the image. One could segment the image, and then determine the area or average diameter of the blanched region at each time interval. One can then declare that the capillaries are filled when the area/diameter of the remaining blanched region is close enough to zero. Alternatively one could segment a region of interest (ROI) from the image and then calculate some average parameter based on the colour of the ROI and log that with respect to time. One can then determine how close that parameter should be to its final value to decide when the capillaries are refilled.

Both of the methods discussed above would be valid ways of determining capillary refill time, unfortunately neither can easily be implemented. The problem with implementing either method is that the data do not lend themselves to being segmented reliably. The difference between blanched and un-blanched skin is faint at best and in some cases almost indistinguishable. It would be possible to write a program that automatically adjusts the segmentation parameters to accommodate varying skin types and lighting. However, it is the

opinion of the author that such an approach will require a disproportionate amount of effort and that a simpler solution should be opted for.

One simpler way of extracting capillary refill time from the data would be to manually define a constant region of interest on the first frame. One can then treat some average colour parameter of the entire region as the indicator for capillary refill. This is the approach that was initially implemented in this project to extract the colour intensity parameters. The extracted parameters are then plotted with respect to time (as shown in 4.10) and the user is asked to select the appropriate CRT from the graph. This method requires a fair amount of user input, however it puts the user in contact with intuitive representations of the data.

Figure 4.3 provides a quick overview of the work-flow for the method of extracting CRT from the raw data, that is implemented in this project.

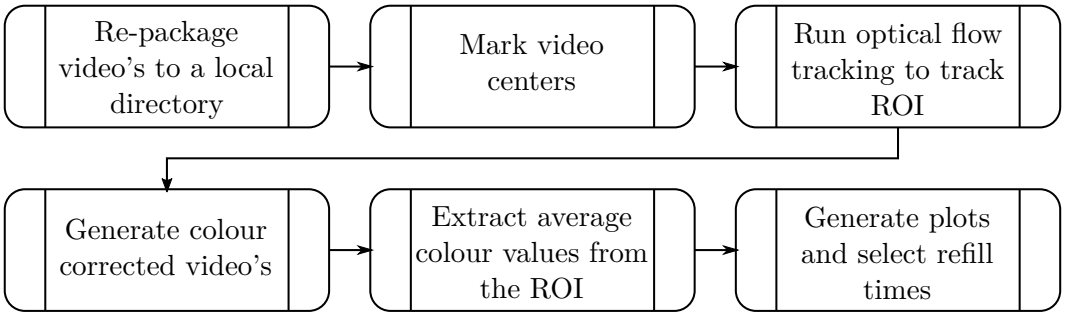

Figure 4.3: CRT analysis work-flow

4.1.1 Repackaging .h264 files

The first step in the process is to copy the .h264 files that were created by the sensor, repackaging them into a .mp4 container and then to save them in a local directory where the application can easily access the data.

As soon as the repackaging process is started, the application builds a dictionary of all the files in the source directory. It then loops through every file and repackages it using FFmpeg. FFmpeg is a GNU licensed set of multi media handling libraries that can be accessed using the windows command terminal and subsequently, python. Repackaging from .h264 to .mp4 with ffmpeg does not involve decoding or re-encoding the data, so no information is lost and the process takes less than a second per video.

This is an important step because the application has uninhibeted write access to the working directory and one programming error in the development phase could destroy all the data.

4.1.2 Marking centres

The next step in the CRT extraction process involves manually marking the blanching site. Figure 4.4 shows what the marking interface looks like. The frame on the left is the 10th frame (0.3 *s* after the camera starts recording) of the CRT video. For a good sample the blanching site should be clearly visible, however this is not always the case. The frame on the right is the exaggerated blanching image which is created to aid in the locating of the blanching site. The exaggerated image is created by taking the difference between the first frame and the last frame in the video and adding a scaled version of the difference matrix to the first frame. Also visable on both images, is a rough hand drawn circle that was used to identify the blacnhing site in the adult study.

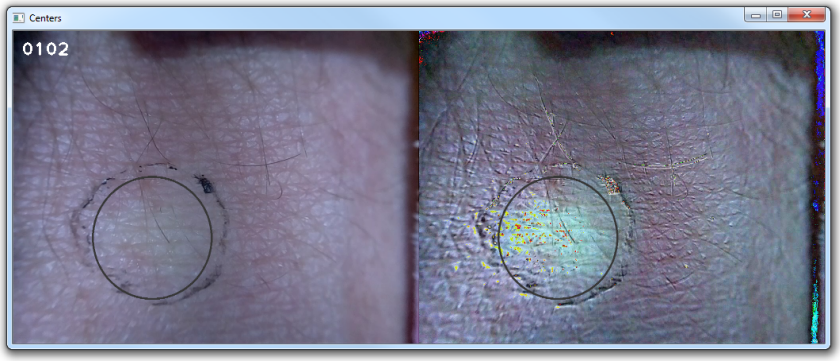

Figure 4.4: Center marking interface for CRT extraction which shows the normal frame on the left and the exagerated frame on the right as well as the selected region of interest

The user can then select the center of the blanching site by clicking on it, on either the original or the exaggerated image. The blanching area is approximated as the circle shown in Figure 4.4. The radius of the circle can be edited by using the '+' and '-' keys. Once the user is satisfied with the placement of the center marker, a simple keystroke saves the coordinates and radius of the blanching site for that video. The two center selection frames for the next video are then automatically loaded. When the process exits the centres are saved as a dictionary to be accessed later. The center coordinates and radii are also displayed on the user interface as shown in Figure 4.5. If the center marking application is started again it will first load the coordinates and radii stored in the previous session and use them as default values for each video. This allows the user to come back at a later stage to make adjustments to one particular set of coordinates without having to re-select all the other centres.

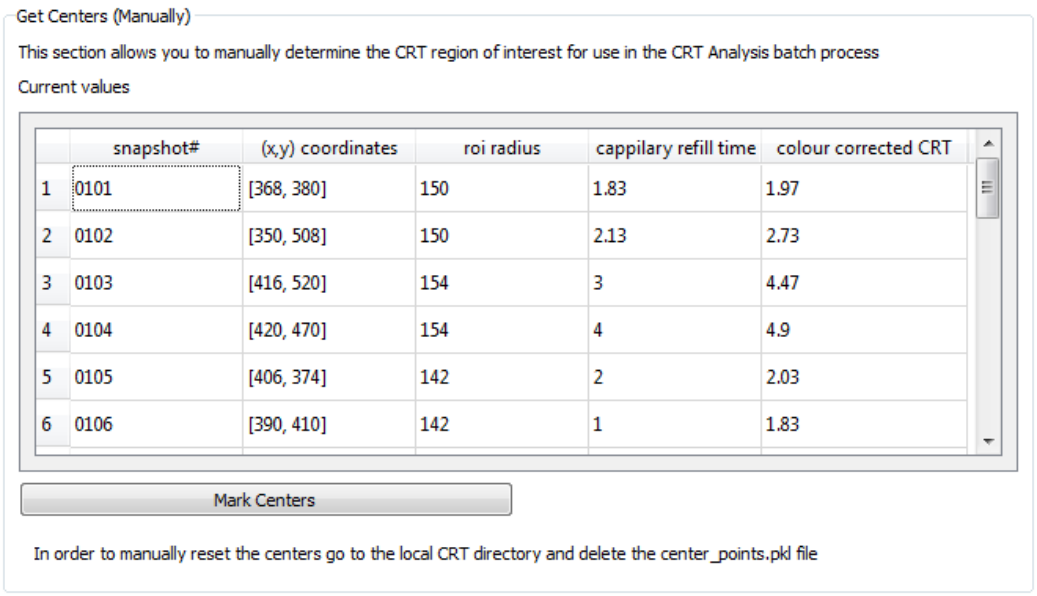

Figure 4.5: Center marker launching interface

4.1.3 Optical flow tracking for image stabilization

The patients that participated in the Tygerberg study were far more restless than originally anticipated. A review of the captured data showed that motion artefacts dominated many of the measurements and that compensation for the motion would be required to maximize the amount of good information that can be extracted from the data. The task of compensating for the motion artefacts can be broken down into two distinct problems. The first is to capture the motion information and the second is to adapt the information extraction process using the captured motion information.

Capturing motion information from the CRT data

Since the need for motion estimation was discovered during the trial, it was not possible to consider the inclusion of a gyroscope or accelerometer in the original design, fortunately, however, there are many ways of extracting the relative motion between the camera and the scene that use only the information in the captured videos. These methods can broadly be categorised as either dense or sparse optical flow estimation. The decision of which flow tracking scheme to use, is based on the properties of the footage that one can exploit, the available hardware and some trial and error.

The CRT sensor data presents a fairly uniform, but textured scene, this means that, with the appropriate prepossessing one can force the feature detection algorithm to find corners in almost any part of the scene. These corners usually correspond to small wrinkles in the patient's skin, as shown in Figure 4.6 below. To limit the computational burden only the 200 'best'

features were tracked. Very generally speaking a good feature in an image is an area in the image that has a large variation in intensity in all directions. Notice how the 'best' features seem to be arraigned in a circle around the center of the image. This is because the skin in this area is at the optimal distance from the lens and is therefore sharper, the area in the middle of the scene is a bit to close and the area outside the circle is a bit to far. The Shi-Tomasi feature detector from the OpenCV library was used to find good features to track. For tracking the points, the sparse iterative version of the Lucas-Kanade optical flow tracking method using image pyramids was used.

Figure 4.6: Frame from a CRT video with the features that are being tracked and the motion traces super imposed on the image.

Very briefly: The Lucas-Kanade method is a method for solving u (displacement in the x direction) and v (displacement in the y direction) for the optical flow equation. Consider a pixel I(x,y,t) in the first frame. If we assume the pixel intensity to be constant between the first and second frame we can write:

$$I(x, y, t) = I(x + \Delta x, y + \Delta y, t + \Delta t) \tag{4.1}$$

Assuming the movement to be small, taking the Taylor expansion of the right hand side yields:

$$I(x, y, t) = I(x, y, t) + \frac{\partial I}{\partial x}\Delta x + \frac{\partial I}{\partial y}\Delta y + \frac{\partial I}{\partial t}\Delta t \tag{4.2}$$

Removing the common terms and dividing by Δt then gives that:

$$I_x v + I_y u + I_t = 0, \tag{4.3}$$

where, $I_x = \frac{\partial I}{\partial x}, I_y = \frac{\partial I}{\partial y}, I_t = \frac{\partial I}{\partial t}, u = \frac{\Delta x}{\Delta t}, v = \frac{\Delta y}{\Delta t}$. This produces an equation with two unknowns (u and v) that cannot be solved directly.

More equations can be obtained by assuming that the neighbouring pixels have the same motion and therefore the same u and v. The Lucas Kanade method (as implemented in OpenCV) takes a 3x3 patch around the point and assumes that all 9 neighbouring pixels have the same motion [20]. The problem then changes to solving 9 equations with two unknowns, which is over determined. The Lucas Kanade method then approximates a solution using the least square error approach. The vectors u and v are then used to approximate the position of the features in the next frame. The feature set (corners) in the scene is very dens, which sometimes means that some of the corners are incorrectly matched which causes the markers that track points on the skin to drift along skin-folds during periods of increased motion. This 'drift' is best illustrated in Figure 4.7 below, which depicts the x and y motion of the 20 best features that were tracked.

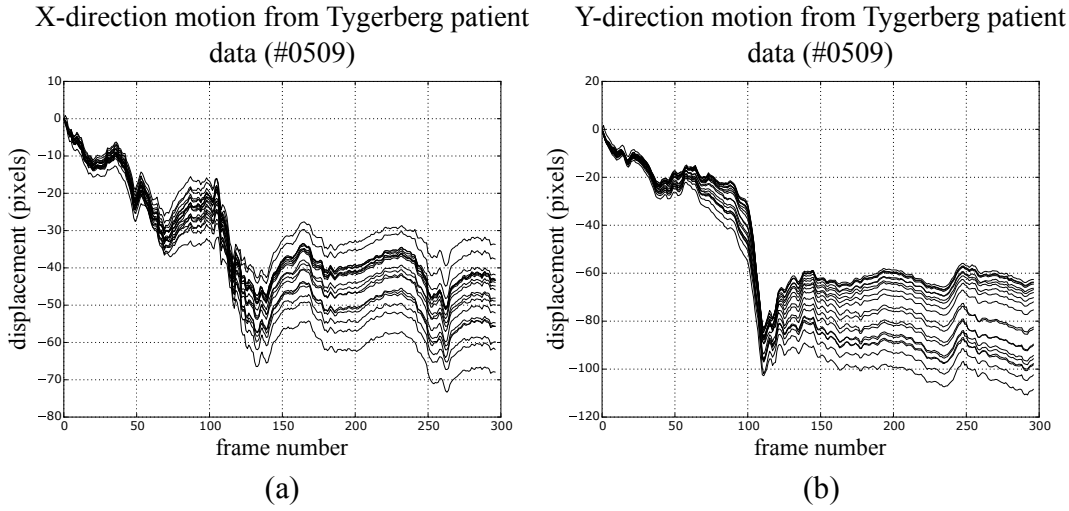

Figure 4.7: X and Y motion of the CRT video of snapshot #0509 from the Tygerberg study data set.

It can be seen that the displacements recorded by each marker diverge slightly. This could be the result of rotation of the image, but manual inspection of the raw data has shown that this is not the case. There are no obvious outliers present in Figure 4.7 so the final motion was calculated by taking the average motion of the best markers. The final global movement vectors were obtained by dividing each frame into four quadrants, calculating the average movement in each quadrant and taking the average of the movement of the four quadrants.

Adjustment of image extraction algorithm using movement vectors

Many ways of adapting the image extraction were considered. Translating and reshaping the image using affine transformations to keep the intended ROI in

the center of the scene seemed promising, however it required an additional conditioning. Instead the ROI extraction algorithm was adjusted to move the ROI that was selected in Section 4.1.2 with the skin.

The use of motion tracking during CRT extraction allowed for the inclusion of an additional 13 samples in the final analysis (an improvement of roughly 60%). Despite this improvement, there were still a significant amount of samples that it could not be included. The remaining samples were either excluded because they had very large movements between subsequent frames that caused the tracking to fail, or the scene moved too far: Every effort was made to provide a uniformly lit scene, however the corners of the scene directly under the LED's were still slightly brighter and in cases where the blanching was not very pronounced it was not possible to distinguish whether the colour change was the result of a change in lighting or the result of capillary refill. Consider the case in Figure 4.8 below. The image on the left represents the scene with exaggerated lighting gradients. The path on the image represents the movement of the ROI as determined by the optical-flow tracking algorithm. On the right one can see the expected colour change due to capillary refill and motion as well as the superposition of the two signals that is returned by the sensor for this case.

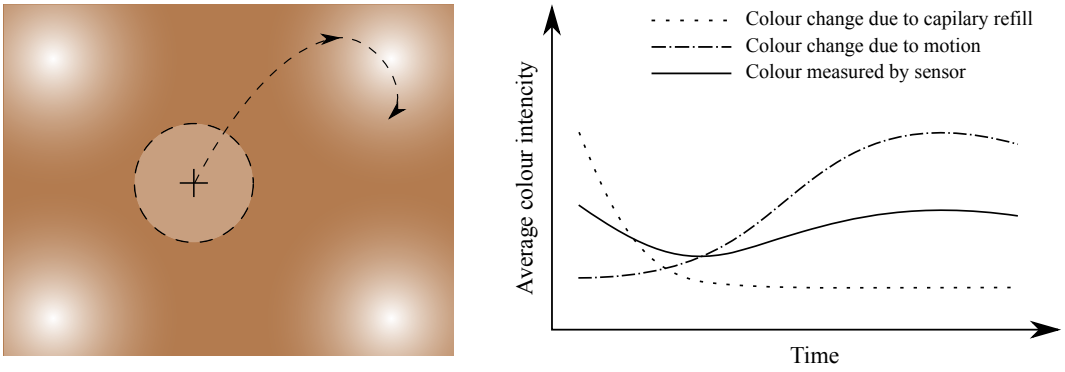

Figure 4.8: Representation of the effect of large motion on the CRT colour data due to non uniform lighting. The motion throught the scene is depicted on the left and the resulting data is depicted on the right.

4.1.4 Colour information extraction and brute force colour correction.

Once the locations of the regions of interest are stored for all the videos it is possible to extract the colour information in the regions of interest for each frame. This project will consider the average colour components in the RGB, HSV and grey-scale colour-space.

The first challenge is to separate the region of interest from the background. This can be done by looping through pixels, but this approach would be very inefficient. Instead the problem will be manipulated such that it can be treated as a matrix algebra problem which will allow us to leverage the efficiency of the numpy matrix manipulation library for Python as well as some features from the OpenCV library. This is done by first creating a binary image mask with the same shape as the frames in the CRT videos. In OpenCV, images are saved as matrices with size: pixel-height x pixel-width x number-of-colour-descriptors. The original images are encoded as RGB pictures which is essentially a 1024x764x3 matrix. This is equivalent to having a 1024x764 matrix where each entry is a 3 digit vector. We then build a mask with ones in the region of interest and zero's outside the region of interest. The background is then removed using the binary operation; $M_{foreground} = M_{original} \& M_{mask}$. With the foreground isolated we now calculate the averages of the non-null pixels in each colour channel. This information is then stored, and the analysis of the following frame begins.

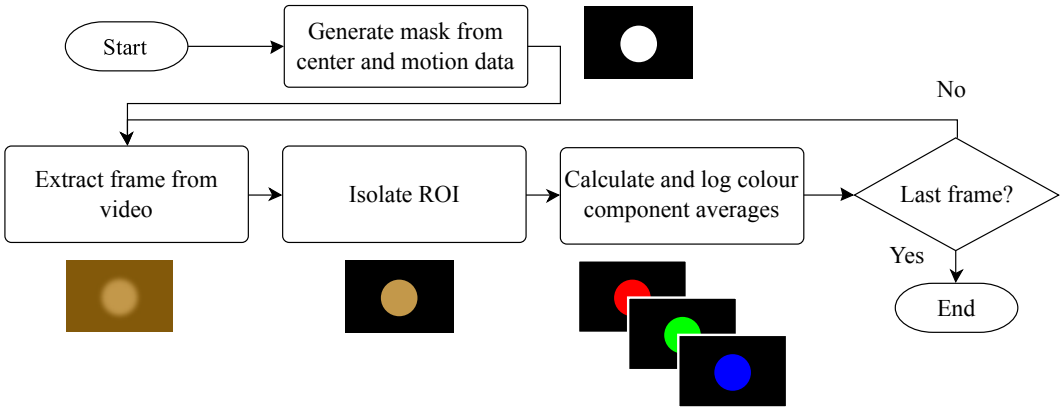

Figure 4.9: Capillary refill colour extraction

The average colour component values are then stored in a Python data-structure for later analysis.

Once the colour information has been extracted for all the CRT videos the user can start the CRT picker interface. This starts an interactive interface that plots the extracted colour values with respect to time as shown in Figure 4.10. The axis labelled `"fg_H"`, `"fg_S"` and `"fg_V"` represent the average foreground hue, saturation and lightness respectively. The axes labelled `"fg_R"`, `"fg_G"` and `"fg_B"` represent the foreground red, green and blue components respectively. `"fg_GRAY"` represents the average foreground intensity when converted to gray-scale and the graph in the bottom right corner depicts the average background colour components in the RGB and gray-scale colour spaces. The x-axis on each of the graphs is frame number. The videos were recorded at 30 fps so the time interval shown is 10 seconds in each case. The dots on

the graphs represent the points that the user has selected as the point where blanching is complete. The user selects these points by simply clicking on one of the graphs. To support in the selection of the correct capillary refill point, a small window opens alongside the CRT picking plots that displays the video frame at the selected time interval. This frame then updates as the user selects a different time on the graphs. Once the user is satisfied with the selected capillary refill time, a keystroke will save the selected capillary refill time and launch the selection graphs and video frame for the next CRT video.

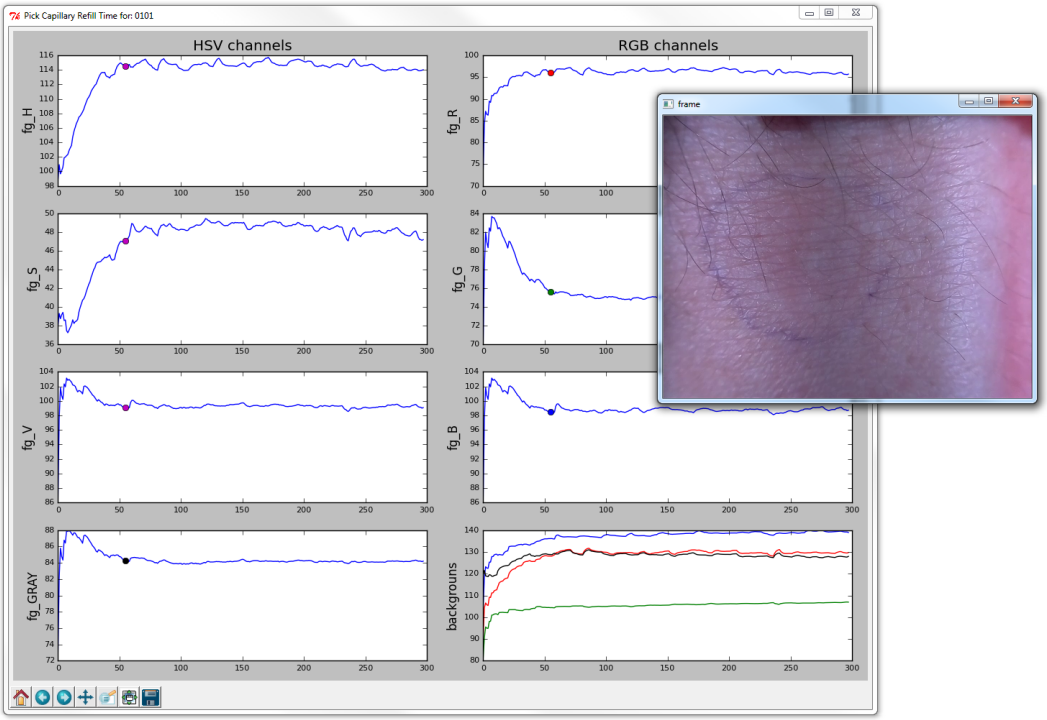

Figure 4.10: Screenchot of the CRT selection interface before brute force colour correction was implemented (colour component intensity plotted against frame number at 30fps) shown with the real-time frame refrince

In the RGB and grey-scale colour spaces, as implemented in Python and OpenCV, each colour component of each pixel can assume a value between 0 and 255, where a higher number indicates a greater presence of the colour and therefore a lighter image. We therefore expect that the average foreground colour components should decrease in the first few frames and then level out. We also expect that the background colour should remain roughly the same throughout the video. In the raw videos we see that this is not the case though. When we look at the average background colour it seems that the camera is going through some transient period, despite the fact that all the automatic colour correction has been turned off. These transients could be caused by some lower level colour correction or they could be the result of some physical

phenomenon like heating of the camera and illumination system. Regardless of the origins of the video transients, it is possible to remove them after recording. The DEA includes a brute-force colour correction script. It is built using the assumption that the background colour should be constant. When the script starts it extracts the average background colour components of the last frame in the CRT video. It then loops through every frame in the CRT video, and calculates the difference between the average background colour of the current frame and last frame and adds that difference to the entire frame. These frames are then re-encoded as a new video and stored. When these videos are then analysed and their colour content plotted as discussed above we obtain graphs as shown in Figure 4.11. Note that Figure 4.10 and Figure 4.11 were generated using the exact same footage. Also visible in Figure 4.11 are the least square approximations for each channel. Before launching the CRT Picker interface the user can select the order of the polynomial to be fitted to the data. At the time of writing, no automated CRT cut-off was implemented, however the fitted curves should be well suited for such operations.

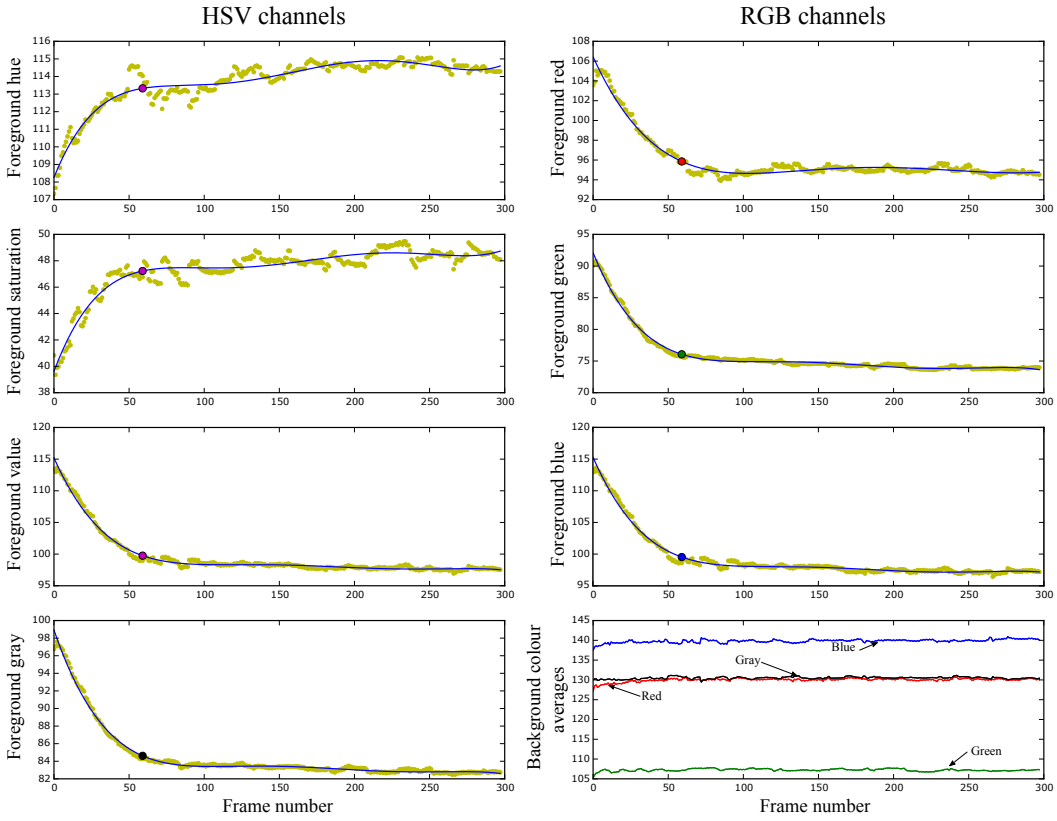

Figure 4.11: Data on CRT picker interface and 4th order least squares approximation after brute-force colour correction

When the user chooses to exit the CRT picker interface or if all the videos have been marked, the selected capillary refill times are saved and also displayed

in the user interface as shown in Figure 4.5. When the CRT picker interface is started again the times selected in the last session will appear as the default values for each test case so that the CRT times may be marked in several sessions. This is useful if one has a particularly large data set.

Examination of all of the Coetzenburg data showed that the intensity in the corrected gray-scale image consistently provided the highest signal to noise ratio.

4.2 Skin recoil time extraction

When examining the raw skin recoil footage closely one will see that, despite the high frame rate and short exposure, the first few frames after the skin is released are quite blurry. This makes any automated information extraction very difficult. A manual skin recoil marking algorithm was therefore implemented.

As with the capillary refill time extraction, the first step is to repackage the raw footage to an easily accessible location in a format that is compatible with OpenCV.

One can then open the skin recoil marking interface from the main user interface as shown in Figure 4.12. The marking interface allows one to fast forward and rewind in increments of 1,10 and 100 frames. The user can then mark the start and stop time of every recoil occurrence using the keys shown in Figure 4.13.

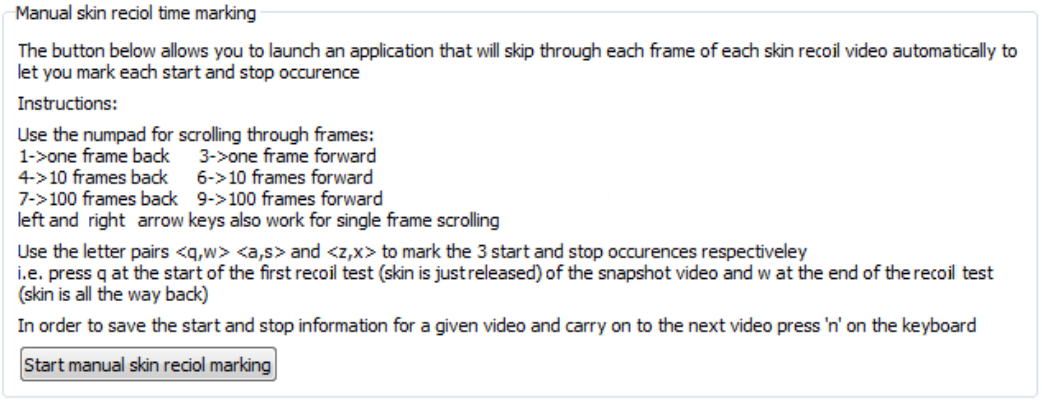

Figure 4.12: Skin recoil section of the main DEA user interface

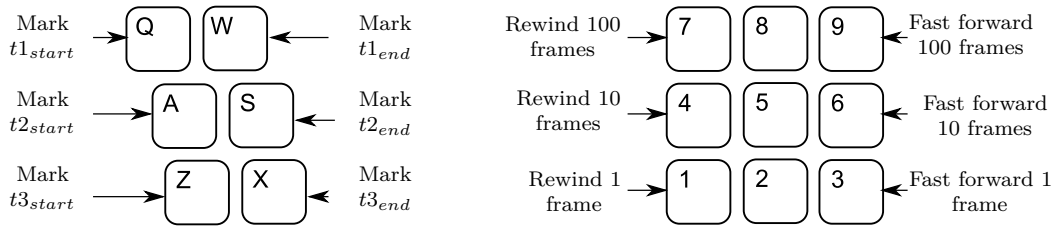

Figure 4.13: Keys for marking skin recoil time using the skin recoil marking interface

Despite the large amount of user input the interface allows for relatively quick processing of the data. The average time for marking one skin recoil occurrence of the Coetzenburg dataset was 20 seconds. As with the other interfaces that require a large amount of user input, the values that are selected in one session are stored and used as the default values of the following session. Figure 4.14 shows four frames that were extracted from early validation tests.

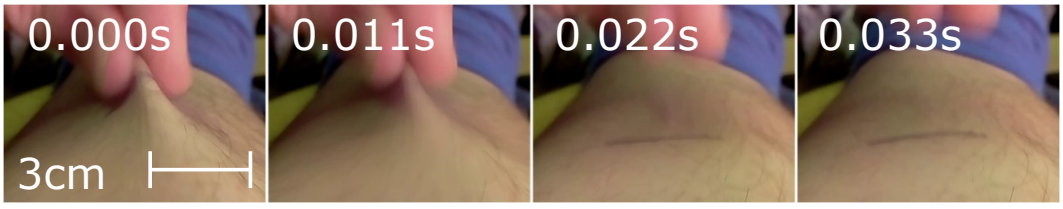

Figure 4.14: An example of skin recoil footage recorded at 90 fps

The frames in Figure 4.14 show a recoil event that was completed in around 0.033 seconds. Recoils that happen at this rate are difficult for physicians to estimate, even to the correct order and would be nearly indistinguishable from a recoil time of 0.066 seconds, for example. The method that is implemented here is only limited by the frame rate of the camera. If this method for measuring SRT is ever to be used in clinical practice it will be necessary to automate the extraction of the SRT. Since this project is only aimed at proving the basic feasibility of the concepts, the method discussed here is sufficient, however for the sake of potential future projects a brief discussion of some ways to automate the SRT extraction process is presented in Section 7.2.

4.3 Temperature profile data extraction

As discussed in the literature review and concept development section, the thermal images will be used to do a proof of concept analysis on a device that can measure a patient's skin temperature along a chosen path. FLIR provide some analysis tools for their product, however these tools are not well suited

for batch processing and therefore a module was added to the DEA application that could extract the required parameters.

Before discussing the thermal profile extraction module, we first consider the raw data. Figure 4.15 shows two examples of the images produced by the FLIR E60 camera that was used in the adult dehydration studies.

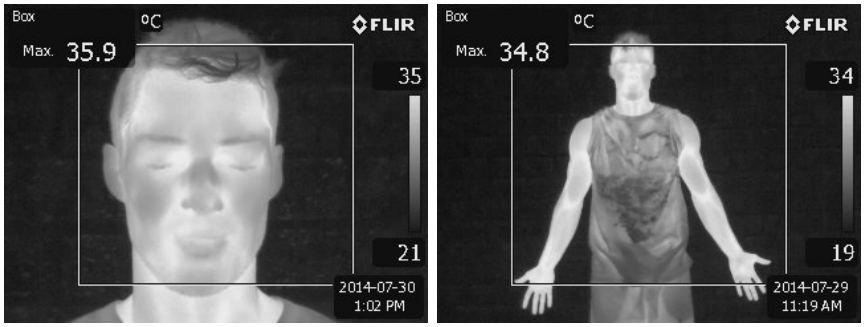

Figure 4.15: Thermal images of a Coetzenburg study participant

The skin surface temperatures are encoded into the pixel values of the image. The camera software automatically calibrates the colours in the output image so that the full range of the pixel values is used when displaying the image. This means that the relationship between pixel intensity and surface temperature may vary from image to image, which necessitates the inclusion of a calibration step in the temperature profile extraction algorithm. The first step in the process involves automatically repackaging the thermal images (.jpg) to match the file structure and labelling convention that has been established for the other sensors using the information from the patient data sheets and the image file names. This can be seen in the first cells of Figure 4.16 which provides a high-level overview of the processes that are involved with extracting the required hydration markers.

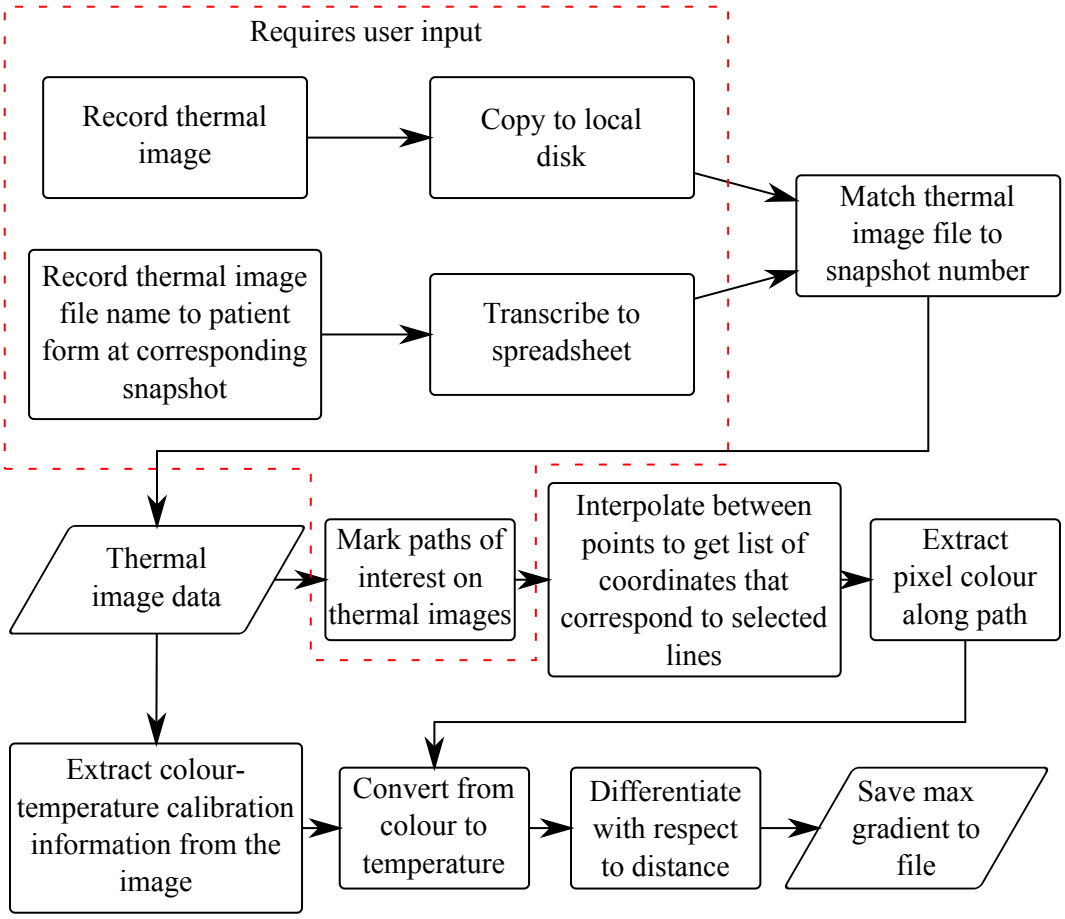

Figure 4.16: Thermal gradients extraction workflow

4.3.1 Thermal and colour calibration

As mentioned above the camera colour pallet adapts to the content of the scene, which necessitates individual calibration of each image. Fortunately all the information that is required to perform the calibration is embedded in the image in the form of the colour-bar and maximum and minimum temperature, as can be seen in Figure 4.15. The first step for the interpolation is to ensure that the colour bar maps linearly when the grey scale intensity is used. To check this the average intensity of each row of pixels of the colour bar is taken and plotted with respect to its y-index, and indeed this results in a near perfect straight line. The maximum and minimum colour intensities are now known as well as the temperatures that correspond to them. The calibration scheme is then set-up as shown in Figure 4.17 where we use the known extreme temperature and colour values to solve the straight line equation. To manually read and note all the temperature values from the images is cumbersome, boring and makes the DEA more susceptible to human error. Therefore as part of the calibration algorithm a subroutine was implemented that extracted the maximum and

minimum temperatures from the image using the Tesseract optical character recognition engine (OCR) that was developed by Google. The Tesseract OCR was accessible to Python through a wrapper function that only requires an image object to be passed to it and returns the text in the image. Extracting the region of the image that contained the required text is trivial since it is always in the same location.

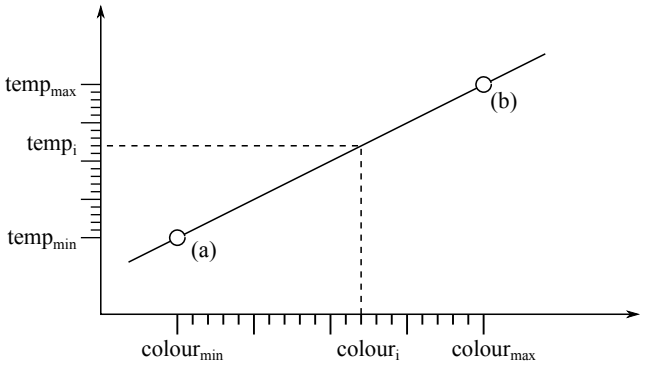

Figure 4.17: Temperature calibration scheme

4.3.2 Marking Paths

The DEA provides a path marking interface that allows the investigator to define the path along which the temperature is to be extracted, by simply clicking on the image. The path is defined as n-1 straight lines between the n points selected by the investigator. If the investigator is satisfied with the selection he can press n to move to the next image or c to clear the selected path and start over. The information is handled in a sensible way so that the process can easily be interrupted and resumed in the case of large data sets.

4.3.3 Path interpolation algorithm development.

Since the hydration marker will be extracted from the gradient of the temperature along the specified path, it is important to ensure that there is no discontinuity in distance between the extracted pixels. It is also important to keep the distance increment constant between line sections as well as patients irrespective of the length of the selected temperature path. This is of course, based on the assumption that the images from the clinical studies were all taken from the same distance, which they were. Consider the 4x3 pixel image in Figure 4.18 below. The empty dots represent the centres of the pixels in the image and the filled dots represent the real coordinates of the interpolated points at which we would like to extract the temperatures.

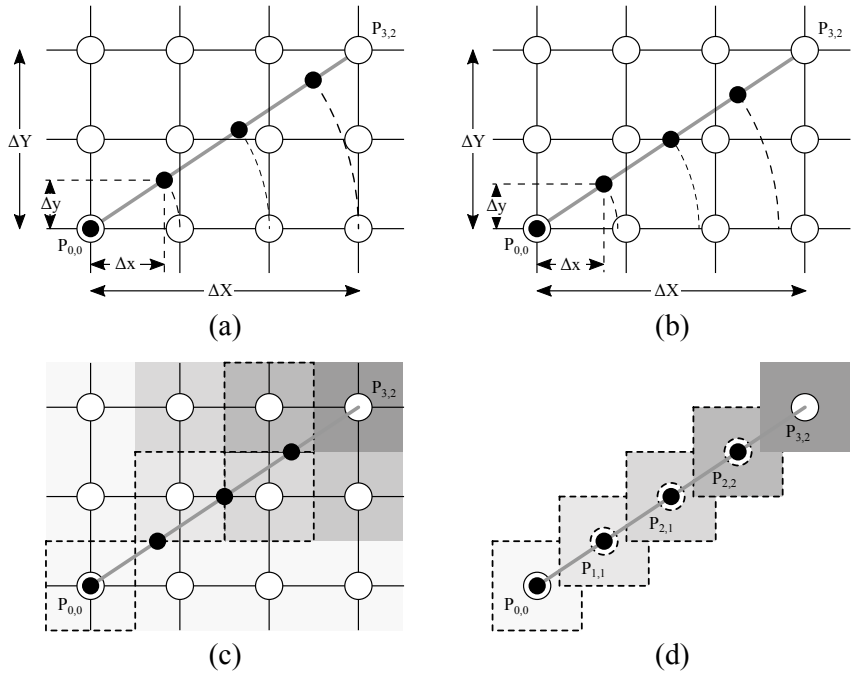

Figure 4.18: Pixel interpolation. (a) Using a one pixel step size. (b and c) Using a step size of aproximately one pixel to eliminate discontnuities. (d) Approximated temperature path

For the path interpolation one pixel was selected as the basic unit length. Assume we would like to get the temperature along the path $P_{0,0} : P_{3,2}$. From the images we see that the length of a path between two points is $dist = |dX + dY|$. One would then be tempted to define the x-increment, dx for the section such that $|dx + dy| = 1$ by calculating dx as $dx = \frac{dX}{dist}$. However Figure 4.18 (a) shows us that this would lead to trouble at the end of the line section, since the diagonal length is almost never an exact multiple of the unit length (of one pixel). One can avoid this situation by breaking up the path into k equal pieces (of approximately one pixel each), where $k \approx |dX + dY|$ and define dx as $dx = \frac{dX}{round(dist)}$ as shown in Figure 4.18(b). In the case of the path in Figure 4.18, it would mean that $dx = 3/4$ and our unit length for this example is then 0.901 in stead of 1. The error in unit length is given by:

$$error = \left| 1 - \frac{|xd + dy|}{round(|xd + dy|)} \right| \tag{4.4}$$

In the above example we get an error of 9.86% which is rather large, however this error quickly decays as the line length increases as shown in Figure 4.19. With the x values chosen, we solve the equation for the line using the two points that define it and calculate the y values that correspond to each x value. The x and y values now represent the ideal location on the image at which we would like to extract the temperature information as shown in Figure 4.18 b

and c. The ideal coordinates are then rounded up to the nearest integer which represents the index for the pixels of interest along the path. Note that the last pixel on the defined segment is not included as it will be added as the first pixel in the next segment assuming that there is one.

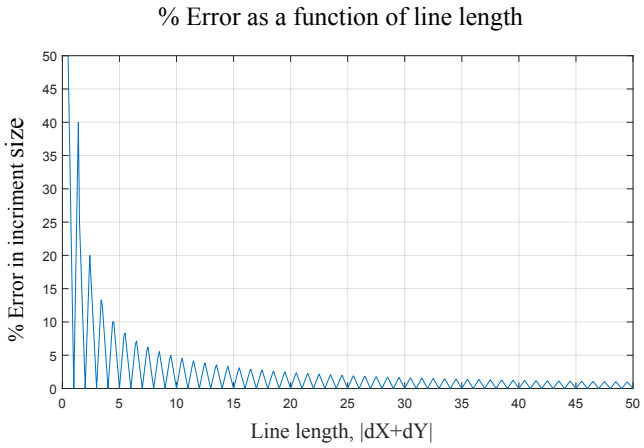

Figure 4.19: Length error as a function of line length

4.3.4 Marker extraction

Using the interpolated pixel path and calibration constants, the temperature profile along the specified path is extracted. The temperature along path is approximated using a 10th order polynomial. The maximum derivative of the approximated temperatures, which is the sought after hydration marker, is then extracted and saved.

5. Study Design

Obtaining clinical data is often difficult and expensive. It is therefore important to have a good system in place to manage the data. For the study all the tests need to be related back to the patient and the specific test index, while anonymizing the data. To do this each patient was given a two-digit ID. Each test event or snapshot was given a four digit ID of which the first two denoted the patient number and the last two denoted the test index. This test ID would then be embedded into all the patient data as it is recorded to minimize the risk of incorrect data labelling. An example of a clinical observation and general observation form can be seen in Appendix E. The QR code for tagging the snapshot can also bee seen on the general observation form.

5.1 Coetzenburg study

As stated in the introduction this project ran parallel with another Masters project at Stellenbosch University which was also concerned with developing non-invasive infant hydration sensors. The Coetzenburg study was developed by the investigator of the other project, however a brief overview of the parts that relate to this project is still appropriate since the devices that were developed in this project were also tested in the Coetzenburg study.

The Coetzenburg study was conducted on 9 adult volunteers (5 male, 4 female, age $= 23.4 \pm 1.59$, BMI $= 24.9 \pm 3.3$), that were dehydrated in one day, by exercise and fluid restriction. Each participant would start the day hydrated and rested. She or he would then be subjected to 3 exercise sessions of 50 minutes each, while not eating or drinking anything. After each exercise session the participants were asked to rest for 20 minutes and the following markers were captured:

- Weight.
- Digital CRT (back of hand).
- Digital SRT (back of hand).
- Thermal image (face and arm).
- Nurse CRT (back of hand).
- Nurse SRT (back of hand).
- Date and time.

This information combined with the ambient temperature as well as the snapshot ID for the test interval would constitute one snapshot of the participant's condition. After the day's exercise and tests the participants were fed, rehydrated and dismissed. This was repeated three times over three consecutive days. Prior to the dehydration tests the patients were asked to adhere to a prescribed diet and rate of fluid consumption in order to obtain well hydrated baseline values.The study was conducted with the approval of the Health Research Ethics Committee of Tygerberg Hospital and Stellenbosch University.

5.2 Tygerberg study

To test the sensors in a clinical setting, a study, that was approved by the Health research ethics committee (HREC) of Tygerberg Hospital and Stellenbosch University (# S13/10/204, "Quantitative Hydration Sensor Development Infant Testing") was conducted at the Tygerberg Hospital Paediatric Emergency and Ambulatory Unit, from here on referred to as G-Ground. The majority of the patients that are admitted to G-Ground are referred from primary care clinics, while a smaller percentage are directly admitted from home. This is not ideal since some of the patients have received some rehydration therapy prior to admission to G-Ground, however this still represents a fair use case for the sensors.

Patients that were considered for the study needed to be admitted to G-Ground with acute gastroenteritis, and be between the ages of 3 and 36 months. Patients that were malnourished i.e. more than 3 standard deviations away of their ideal weight (as determined by the WHO weight for age charts) were not considered for the study. Patients with chronic conditions related to diarrhoea or perfusion were also not considered for the study. What constituted a chronic condition related to diarrhoea or perfusion was determined on an ad-hoc basis by the attending physician.

As soon as a suitable candidate was identified, informed consent was obtained from the parent or guardian. The informed consent forms were reviewed and approved by the HREC and translated from English to Afrikaans and Xhosa. The English version of the informed consent form can be viewed in Appendix E. As soon as informed parental consent has been obtained, the initial study admission form (also in Appendix E) was completed. This form includes all the information that is required to identify the patient as well as notes on the patient's condition at admission. As soon as the study admission form is completed, the first patient snapshot is taken.

For the Tygerberg study one patient snapshot consists of three parts, namely, the routine assessment, the clinical assessment, and the digital hydration marker measurement. The routine assessment is performed at admission and thereafter

six times a day or at the convenience of the nursing staff. During the routine patient assessment, the nurses measure the patient weight, heart rate, breathing tempo and temperature. They also record the amount of fluids and solids the patient consumed and excreted. The clinical patient assessment is conducted by one of the on duty doctors in the ward at admission and thereafter roughly every four to six hours. During the clinical assessment the doctor assesses the patient's general appearance and provides an assessment of the patient's hydration level based on gestalt. Additionally, for the purpose of the study, the attending physician will also record the patient's skin recoil time, capillary refill time, neurological state, eye- and fontanelle appearance, mucous membrane dryness, tear production and pulse abnormalities. Lastly, the digital hydration measurement is performed by the study investigators, under the supervision of one of the nurses in the ward. In order to minimize the disruption that the patient experiences, the digital sensor measurements were done directly before or after the routine measurements. During the digital hydration marker measurement the patient's skin recoil time, capillary refill time and thermal profile are recorded. The routine assessment and digital hydration marker measurements happened synchronously, so that all the data gathered during this instance would represent one point in time (snapshot). The clinical assessments unfortunately had to be done at the convenience of the on-duty physicians. Time and date stamps were therefore later used to consolidate the clinical marker data set with the others.

The detailed procedure for each of the devices as they were used in the trials are shown below. For the CRT sensor:

Capillary refill time sensor

1. Turn on the sensor (startup takes <1 minute).
2. Load QR code from routine check-up form (<2seconds).
3. Disinfect hands and sensor contact surfaces.
4. Using a medical marker, make a small dot on the sternum of the patient (if not already present)
5. Lock the arm into position, and with the patient lying on his/her back place the sensor over the sternum so that the pressure pad touches the marked position.
6. Apply light downward force to the sensor. As soon as an applied force is detected it is displayed on the screen in the form of a bar on the left-hand side of the display. The operator should take care to ensure that the correct force is applied by ensuring that the bar remains on the 5th (yellow) pixel of the display.
7. The application of the downward force also triggers a timer that will start counting down from 6 seconds and is displayed as a bar on the right of the display.

8. As soon as the timer runs out, pull the lever that allows the spring-loaded arm to swing out of the way so that the camera can start recording.
9. As soon as the device starts recording the lighting LED's turn on. The recording will run for 10 seconds. Once the recording is done the LEDs turn off and it is safe to remove the sensor.

Skin recoil time sensor

1. Turn on the sensor (startup takes <1 minute).
2. Load QR code from routine check-up form(<2seconds).
3. Disinfect hands and sensor contact surfaces.
4. Place the sensor on the sternum of the patient with the movable camera head facing the abdomen. During placement ensure that the directional LEDs overlap on an accessible part of the abdomen and start the recording.
5. Lightly pinch and pull the skin at the area where the LEDs overlap to mimic the clinical skin recoil time procedure. Do this three times.

Thermal camera

1. With the patient lying on his or her back, take a picture of the patient's face and torso from 60 cm away and at a 45 degree angle. Ensure that both the face and torso are in the frame.
2. With the patient still lying on his or her back, take a picture of the patient's legs and feet from 60 cm at a 45 degree angle. Ensure that both the patient's waist and feet are in the frame.
3. Record the image file names that are displayed on the camera screen on the routine assessment form.

During the sensor operation it is important to keep the patient pacified so that the sensors can be kept as still as possible. It was found that the most effective way of doing so, was by encouraging a family member to engage their attention while the tests are carried out.

The patients were discharged when the physicians decided that they were stable enough. Prior to discharge the patients are weighed one last time. This weight is used as the post-illness weight and is assumed to be the rehydrated weight. In cases where the patients were moved to other wards before being discharged, the last weight recorded by the ward to which they were sent was used as the post illness weight, when possible.

6. Data Analysis and Results

The purpose of this project is to develop a set of sensors that could help to improve the treatment of infant dehydration as a result of diarrhoea, by improving the dehydration assessment accuracy in an underserved clinical setting. Treatment is usually based on the classification of the dehydration level which can be either mild, moderate or severe (>5%,>7%,>10%). Therefore the ability of our sensor to classify the patients according to the above categories would be a good performance measure. Mild and moderate dehydration cases are usually treated similarly, and can therefore be treated as one category. The following analysis will therefore be primarily concerned with the 5 and 10% classification thresholds.

Many of the guidelines that dictate how dehydration should be treated are based on studies that immediately reduce the collected data to categorical values, because of practical limitations in the measurement methods. The predictors that were recorded in this study (as well as the ones that would be recorded in practice if this approach were to be adopted), were recorded as real values, it would therefore be interesting to see how close one could get to a 'proportional' assessment of dehydration, using the recorded predictors. In other words, what confidence limits should we assume in order to report dehydration level as a percentage?

The data analysis will consist of two sections. The first part will only look at the device's ability to record the desired predictors accurately. In the second section we will describe and quantify the relationships between the hydration markers and participant's level of dehydration.

6.1 Device technical performance

6.1.1 CRT sensor

The CRT sensor consists of two separate sensors, namely, the camera and the force sensor that is used to regulate the blanching pressure.

Camera performance

To assess the camera performance we will look at the measurement to noise ratio of the colour signal as well as the 'interpretation uncertainty', which will be defined later. Figure 6.1 below shows a graph of the mean gray-scale intensity over time of the blanching region of a patient from the adult study just after the blanching force has been removed. The dots represent the raw gray-scale values (0=black, 255=white) and the line represents a 4th order fit of the data.

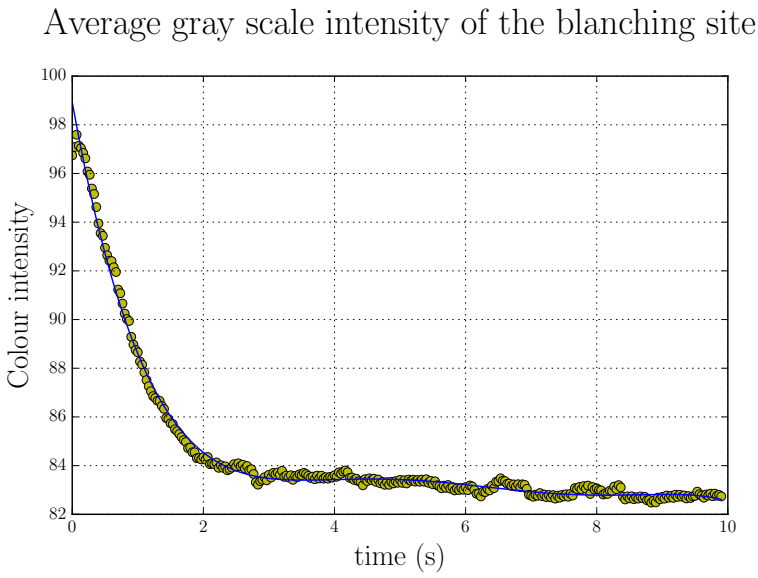

Figure 6.1: Colour intensity at blanching site

We will now define the noise as the range of the difference between the raw data and the approximation, which is clearly an over estimation of the average noise. The measurement size is defined as the range of the approximated signal in the domain of 0-10 seconds. If we then calculate the average measurement to noise ratio of all the tests done during the adult study we get that the average noise contribution of the sensor is roughly 9% of the measurement. A bigger uncertainty however is introduced by the interpretation of the graph. For this study the CRT time was manually determined by looking at the graphs and selecting a point on the curve. This inevitably leads to the inclusion of some operator error. In an attempt to quantify this the operator marked and cleared the CRT times of each snapshot in the adult study dataset three times (3×150 samples). It was found that the average range for each snapshot was around 0.3 seconds or 15% of the average recorded CRT. It is important to note that this uncertainty, which will from here on out be called the interpretation uncertainty, exists because the study investigators opted to determine the thresholds manually. This uncertainty is therefore not an inherent part of the

sensor, but rather a result of the specific way in which it was used in the study. A similar analysis was carried out on the infant study data. Unfortunately there were large motion artefacts in the signals which meant that a measurement to noise assessment as described above would be misleading. However, since the same camera was used it would be fair to assume similar behaviour. A bigger concern for the CRT performance on the infant data was the average interpretation uncertainty, which was hugely exacerbated by the aforementioned motion artefacts. In order to improve the integrity of the CRT data many of the datum points with motion artefacts near the CTR threshold were omitted. In total 15 of the 51 samples were rejected. The average interpretation uncertainty of the remaining samples was approx 35% of each measurement.

Force sensor performance

The Honeywell FSS 1500 has a quoted repeatability of 10 g at 300 g. It is also mounted directly above the pressure pad so it is reasonable to assume that the force measured by the force sensor is the same as the force experienced by the patient at the blanching site. The display was calibrated in such a way that the bar is on the 5th pixel (yellow pixel) in the range of 2.7 N to 3.3 N. During the study the researchers had no problem keeping the force bar at the correct height during the blanching phase of testing.

6.1.2 SRT sensor

The skin recoil time assessment accuracy was limited by the frame rate of the camera, the skin recoil time interpretation uncertainty and most of all the consistency with which the patient's skin was pinched and pulled (actuated). In the SRT sensor, the camera was set to record at 90 fps. Therefore the maximum error that can be attributed to the camera is 5.5 ms. It is difficult to isolate the interpretation from the actuation error, so for this analysis they are lumped together. For both studies the SRT was determined by taking the average of three consecutive skin recoil measurements. It is therefore sensible to define the lumped average sensor error as the average range of the three measurements at each snapshot. This is best illustrated with the graphs in Figure 6.2. The graphs show each of the three test cases plotted against the mean of those cases. This provides an intuitive representation of the measurement to noise ratio as discussed above.

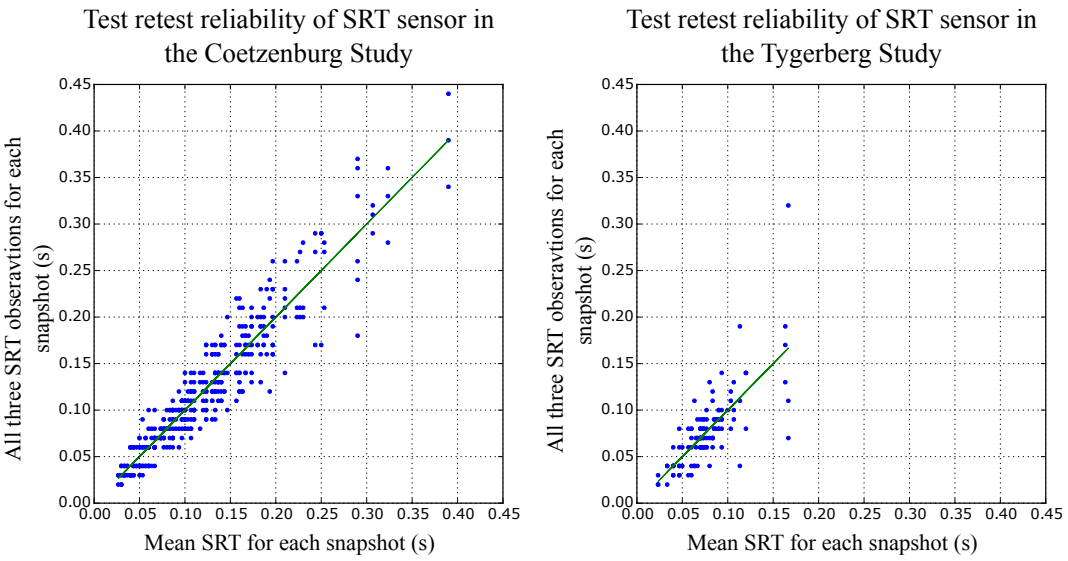

Figure 6.2: The test re-test reliability of the SRT sensor

It would then make sense to quantify measurement uncertainty as:

$$\text{average normalised measurement uncertainty} = \frac{\sum_{all\ snapshots}\left(\frac{range(SRTs)}{mean(SRTs)}\right)}{n_{snapshots}} \tag{6.1}$$

Then, according to the above equation the SRT sensor had a average measurement uncertainty of 42% during the Coetzenburg study and an average measurement uncertainty of 48% during the Tygerberg study. This can visually be verified from Figure 6.2.

6.1.3 Thermal camera sensor performance

A calibrated FLIR E60 was used for the tests. The FLIR E60 datasheet reports a <0.05 °C thermal sensitivity and a thermal accuracy of 2% of the reading, so roughly 0.7 °C . Since the gradients of the thermal profiles that are derived from the thermal images will be used in the hydration assessment, it is also important that the scale of the images are constant. It is fair to assume that the infant's circulation strength can be normalised with respect to their size. In which case the apparent size of the infants in the infrared image need to be the same in order to ensure consistent results. A manual inspection of the infant dataset showed that the largest difference in apparent size of the infants is less than 10%, therefore it is fair to assume that the maximum error as a result of the scale of the image is around 10% of the measurement.

6.2 Clinical results

The clinical performance of the sensors will be based on the results from the Tygerberg study, since this most accurately represents our intended population. During the Tygerberg study a total of 10 (6 male and 4 female) patients were recruited. The median age of the patients was 10.0 months with an interquartile range of 5.0 months. From these 10 patients a total of 51 snapshots were taken. A summary of the study participants can be seen in Table 6.1 below.

Table 6.1: Admission summary for the Tygerberg Patients

Patient Number	Age (months)	Gender	Admission Weight (Kg)	Discharge Weight (Kg)	Percent Dehydration	Length of Stay (days)	Weight z-score
1	9.6	m	10.25	10.42	2.1	0.9	1.265
2	9.3	f	6.46	7.30	11.0	2.6	-1.070
3	10.3	f	6.60	6.97	3.9	1.6	-1.730
4	5.2	f	5.10	5.85	12.8	0.9	-1.463
5	14.0	m	7.50	8.56	12.4	3.9	-1.509
6	12.0	f	9.16	9.39	2.0	1.1	0.356
7	5.7	m	6.14	6.43	-1.1	0.9	-1.770
8	5.7	m	6.05	6.18	2.9	0.9	-2.117
9	10.4	m	9.49	9.48	-0.2	1.7	0.206
10	29.8	m	11.00	12.75	13.7	1.9	-0.357

The last column in Table 6.1 shows the z-score of the patient weight with respect to age at discharge. This denotes how many standard deviations a patient is away from their ideal weight according to the WHO growth charts. This information is repeated in the graphs on Figure 6.3. The lines on Figure 6.3 represent the standard deviations (solid line = ideal weight). In other words, a patient that would fall on the top line for example would be 4 standard deviations above their ideal weight.

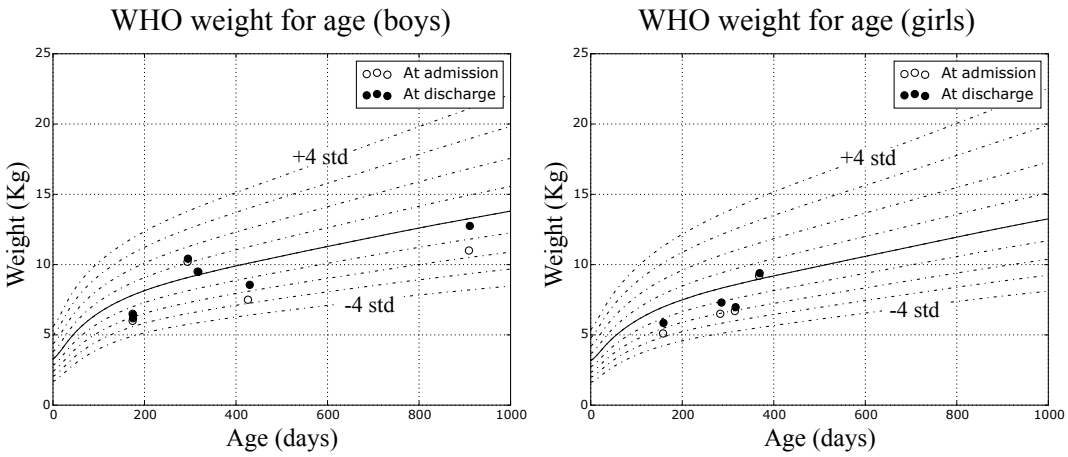

Figure 6.3: Patient weight and age at admission and discharge with WHO growth charts as reference

6.2.1 Reference clinical results

The clinical performance of the sensors will be presented in the following section, their performance will primarily be rated based on their sensitivity, specificity and receiver operating characteristic (ROC) curves. In order to know what constitutes a good sensitivity and specificity and what a good ROC curve would look like we will first consider the performance of the existing methods from literature, we will then consider the performance of the existing methods in our study and then finally, we will present the clinical performance of our sensors. First, however, some review. Consider the two patient populations in Figure 6.4(a).

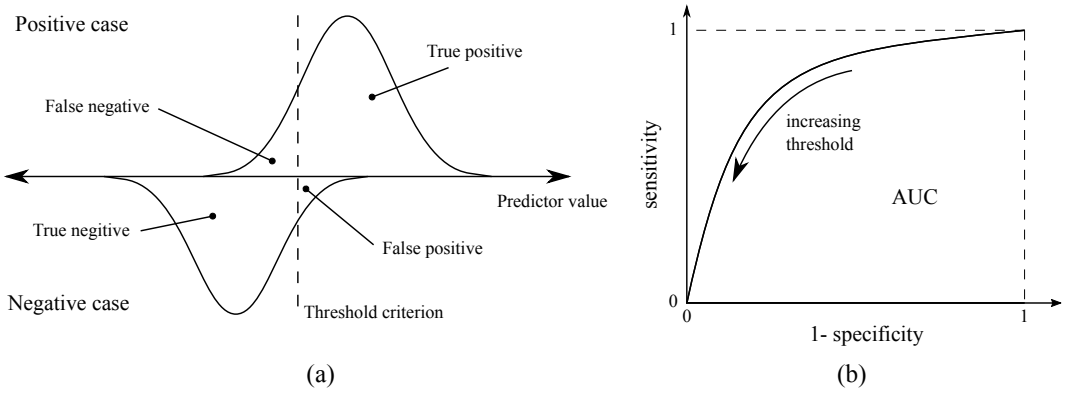

Figure 6.4: Graphical representation of sensitivity, specificity and a general ROC curve

Above the horizontal axis are the positive cases, or the cases that are afflicted with the condition to be identified, in this case, dehydration of more than 5%.

Below the line are the negative cases, the patients that are not dehydrated above 5%. The horizontal axis represents the value of the predictor for the specific condition, for example CRT or SRT. The dashed vertical line represents our threshold criterion. All patients with a predictor value greater than the threshold value are classified as positive and all patients with a predictor value smaller than the threshold are classified as negative. The patients that have the condition and are correctly classified as positive are called the true positive (TP) group, the patients that have the condition and are classified as negative are called the false negative (FN). Similarly, patients that don't have the condition and are correctly classified as negative are called the true negative (TN) group and the patients that don't have the condition and are classified as positive are called false positive (FP).

Sensitivity and specificity are then defined as:

$$\text{sensitivity} = TP/(TP + FN) \tag{6.2}$$

$$\text{specificity} = TN/(TN + FP) \tag{6.3}$$

Or more intuitively; the portion of the population distribution for which the predictor value is on the correct side of the predictor threshold in Figure 6.4(a). The data for the ROC curve is generated by placing the threshold line on the far left of Figure 6.4 and recording the sensitivity and specificity of the sensor classification while incrementally moving the threshold line to the right of the figure. One then plots sensitivity against 1-specificity to create the ROC curve as shown in Figure 6.4 (b). The area under the ROC curve (AUC) is used to describe the performance of a sensor, without imposing the use of a specific threshold value. A larger AUC implies greater sensitivity and specificity for a given threshold.

Reference results for individual indicators

In 2004 Steiner published an article that reviewed all the clinically reliable markers for dehydration in infants [12]. In his conclusion he states that the literature only shows fair to moderate agreement among examiners. The relevant individual clinical markers for the assessment of dehydration are listed in Table 6.2 below for reference. Note the quoted sensitivity and specificity for CRT and SRT from the Tygerberg study, that are presented in Table 6.2, are that of the categorical, manually assessed CRT and SRT.

Table 6.2: Summary of relevant clinical markers for reference, adapted from [12] and original data

Markers	Data source	Sensitivity (95%CI)	Specificity (95%CI)
Prolonged capillary refill	Steiner review	60% (29 to 91)	85% (72 to 98)
	Tygerberg study	30%	73%
Abnormal skin turgor	Steiner review	58% (40 to 75)	76% (59 to 93)
	Tygerberg study	13%	85%
Cool extremity	Steiner review	10% (1 to 40)	89% (80 to 98)

Reference results for clinical scales

As stated in the literature review, the most common clinical hydration scales are the Gorelick, CDS and WHO dehydration scale and between the three of them they represent the gold standard for non invasive, clinical, prospective hydration assessment. The performance of the scales, as claimed when they were first proposed will be presented here as well as the performance of the scales according to two papers that reviewed their performance. The first paper tested the scales in a regional paediatric referral hospital in Rhode Island, USA, while the second tested the scales in three Rwandan hospitals. It is important to note that the data from the Rhode Island study over estimates the performance of each scale since it did not make use of the prescribed threshold values for each scale, but instead used thresholds that provided the optimal performance for the tested cohort.

The data that was collected during the Tygerberg study was also used to reconstruct the hydration scores for each scale. The performance of the hydration scores using the recorded study data is shown in Table 6.3 below.

Table 6.3: Summary of relevant clinical scales for reference, adapted from [21, 14, 13], as well as results from original data.

Dehydration scale	Data source	Sensitivity (95%CI)	Specificity (95%CI)
CDS	Original claim	N/A	N/A
	Rhode Island study	83% (52 to 98)	55% (45 to 65)
	Rwanda study	68%	45%
	Tygerberg study	100%	20%
Gorelick (10)	Original claim	87%	82%
	Rhode Island study	75% (43 to 95)	55% (45 to 65)
	Rwanda study	21%	82%
	Tygerberg study	56%	56%
WHO	Original claim	N/A	N/A
	Rhode Island study	25% (5 to 57)	84% (76 to 91)
	Rwanda study	50%	61%
	Tygerberg study	13%	79%

6.2.2 CRT sensor clinical performance

In the infant study the CRT sensor performed poorly. The digitally measured CRT (dCRT) showed no correlation to the reference hydration level (p= 0.388 R2=0.023) as can be seen in Figure 6.5. In fact, the regression line of the scatter plot between the capillary refill time and the reference hydration level has a slight negative slope, which is the opposite of what one would expect. The ROC curve on the right of Figure 6.5 shows that no threshold for digital CRT would allow the sensor to predict 5% dehydration with acceptable sensitivity and specificity. Adjusting for change in ambient temperature, by performing a linear regression, that related a combination of CRT and ambient temperature to observed dehydration, did not improve the correlation.

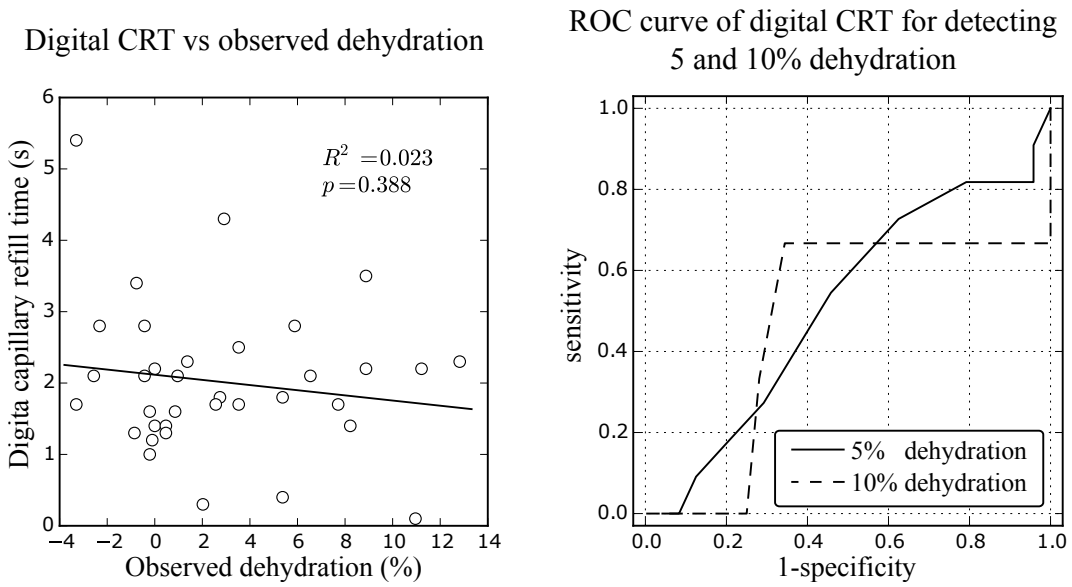

Figure 6.5: Scatter plot and ROC curve for CRT sensor

During the study the physicians were asked to record the CRT as well, using a categorical scale where the categories were: *normal*, *prolonged* and *very prolonged*. Using a *prolonged* CRT as the cut-off for 5% dehydration resulted in an assessment with a sensitivity and specificity of 0.31 and 0.74 respectively. Using *very prolonged* CRT as the cut-off for the assessment of >10% dehydration produces a sensitivity and specificity of 0.43 and 0.98 respectively.

6.2.3 SRT sensor clinical performance

The digitally measured skin recoil time (dSRT) correlated well with the reference hydration level (R2=0.247, p<0.001), as shown in Figure 6.6. The AUC of the ROC curve for the SRT sensor in Figure 6.6 is 0.80. This is very good for

any single non-invasive prospective infant hydration assessment tool. If one were to select a threshold of 0.08 seconds and classify all children in the study with a longer recoil time as being at least 5% dehydrated, one would get a sensor with a sensitivity and specificity of 0.8 and 0.84 respectively for this dataset. A quick glance over Tables 6.2 and 6.3 shows that this combination of sensitivity and specificity is more diagnostic than any single clinical marker that was reviewed by Steiner and is on par with the original performance that was reported by Gorelick for the 10 point hydration scale.

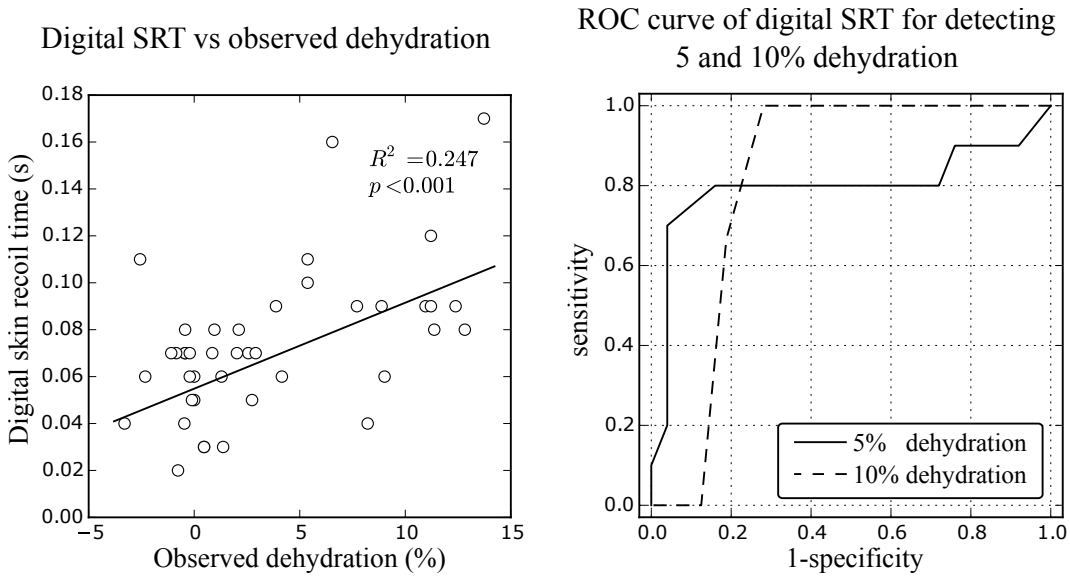

Figure 6.6: Scatter plot and ROC curve for SRT sensor

Despite the good performance that was recorded by the SRT sensor, the SRT as assessed by the clinicians (*instant*, *slightly prolonged*, *significantly prolonged*) appeared to have no diagnostic value with a sensitivity and specificity of 0.13 and 0.85 respectively, when using *slightly prolonged* as the threshold for 5% dehydration.

6.2.4 Thermal profile sensor clinical performance

As discussed in Section 4.3 the marker that would be extracted is the maximum skin surface temperature gradient along a path, specifically along a line on the abdomen for the infant patients, as shown in Figure 6.7.

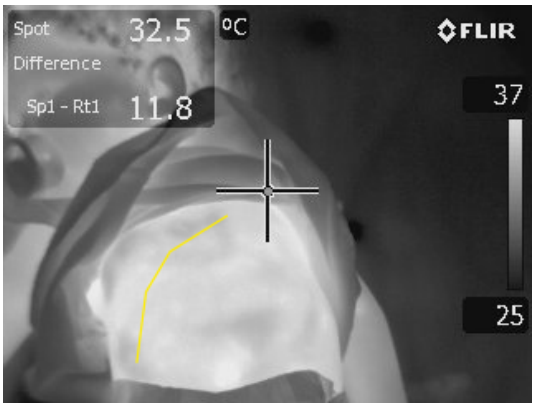

Figure 6.7: The path on the abdomen along which the temperature profile is extracted.

This measurement site was chosen because it crosses areas of both high and low local heat output [22, 23]. It is therefore reasonable to expect larger temperature gradients with lower peripheral perfusion along the specified path. The performance of the thermal profile sensor was fairly mediocre. It performed better than the CRT sensor, but not as well as the SRT sensor. With the appropriate threshold it could achieve a sensitivity and specificity of 0.9 and 0.5 for the infant dataset as can be seen in figure 6.8.

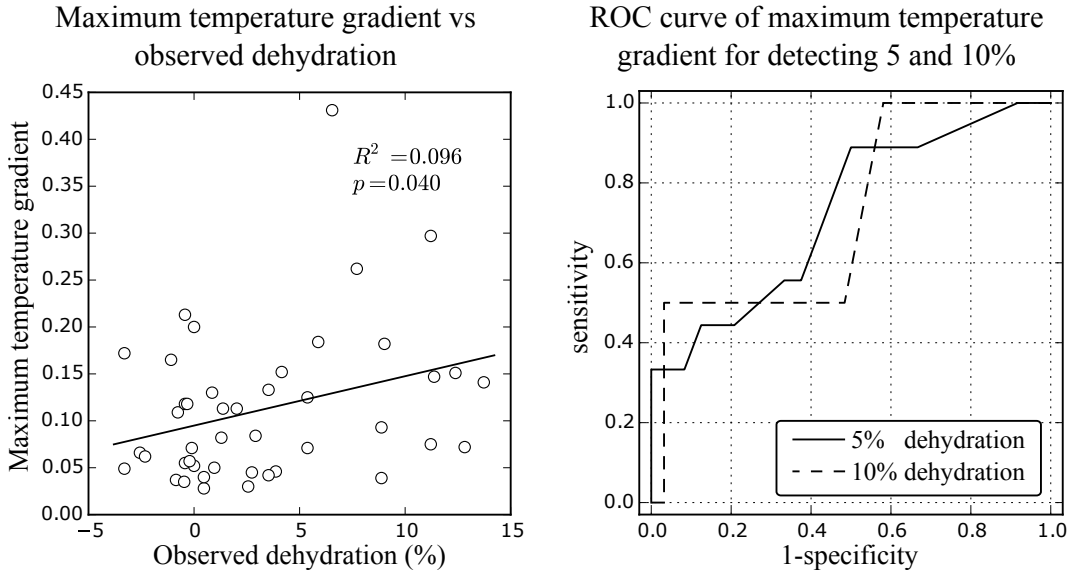

Figure 6.8: Scatter plot and ROC curve for Thermal Profile sensor

6.2.5 Combined predictor performance

There is still no evidence of a single in-vivo measurement that has the ability to classify hydration state sufficiently, it is therefore common in the field of

dehydration assessment research to examine the performance of some model that is based on the combination of a set of dehydration predictors [24, 1, 13]. This is often done, by finding a set of easily measurable predictors with a reasonable degree of orthogonality between them and relating a combination of these measurements to the hydration level (usually determined by post illness weight gain). The data from the Tygerberg study (by design) is well suited for such analysis and since the sensors provide a somewhat novel form of the well known markers, it would be interesting to see if a model could be derived from the available information that would be significantly better at predicting dehydration than any of the input components. This analysis is especially interesting for this sensor configuration, because all the markers will be available on the proposed sensors as soon as they are measured, so this analysis could be carried out during the initial examination without any additional input from the physician.

To start this analysis we will first inspect and visualise the data. The set of predictors that will be combined consists out of the three digitally measured markers as well as age-normalised heart rate and respiratory rate, which are included, since they will most likely be available in the proposed use-case and are known to be linked to hydration state. This selection produces a 5-dimensional set of predictors for each hydration state, which is difficult to visualize. In order to better visualize the data, a principle component analysis (PCA) will be conducted to project the data onto a 2-dimensional subspace. The PCA allows one to project the data from an M dimensional space to an N dimensional subspace, while retaining the most information possible by selecting the subspace in such a way that the variance of the data in the resulting subspace is maximized. If one represents the predictor dataset as a 51x5 matrix where each column represents a predictor(dCRT, dSRT, max temp gradient, normalised respiratory rate, normalised heart rate) and each row an entry. Each sample row can then be depicted as a 5-dimensional vector:

$$x_i^T = \begin{pmatrix} x_{i,1} \\ x_{i,2} \\ x_{i,3} \\ x_{i,4} \\ x_{i,5} \end{pmatrix} = \begin{pmatrix} \text{digital capillary refill time}_i \\ \text{digital skin recoil time}_i \\ \text{maximum temperature gradient}_i \\ \text{age-normalised heart rate}_i \\ \text{age-normalised respiratory rate}_i \end{pmatrix} \tag{6.4}$$

The collected data is in different units and the elements have vastly different magnitudes, so the first step in the analysis is to normalize each column by subtracting the mean of the column from each entry and dividing each entry by the standard deviation (also know as z-transform or normalization). Once the data is normalized and centred we calculate the co-variance matrix. The elements of the covariance matrix are the covariance between two features, which can be calculated as follows.

$$\sigma_{jk} = \frac{1}{N-1}\sum_{i=1}^{N}(x_{ij} - \bar{x_j})(x_{ik} - \bar{x_k}) \tag{6.5}$$

The symmetric covariance matrix $\sum_i$ can then be written as:

$$\sum_i = \begin{bmatrix} \sigma_1^2 & \sigma_{12} & \sigma_{13} \\ \sigma_{21} & \sigma_2^2 & \sigma_{23} \\ \sigma_{31} & \sigma_{32} & \sigma_3^2 \end{bmatrix} \tag{6.6}$$

For the structure as shown in Equation 6.4 and 6.6 the covariance matrix is:

$$\sum_i = \begin{bmatrix} 1.02 & -0.165 & -0.028 & 0.037 & 0.262 \\ -0.165 & 1.02 & 0.33 & 0.205 & -0.042 \\ -0.028 & 0.33 & 1.02 & 0.046 & 0.192 \\ 0.037 & 0.205 & 0.046 & 1.02 & 0.001 \\ 0.262 & -0.042 & 0.192 & 0.001 & 1.02 \end{bmatrix} \tag{6.7}$$

The eigenvectors of the covariance matrix then point in the directions of maximum variance and the corresponding eigenvalues represent the magnitude of the variance in these directions. For our dataset we then have:

$$\text{Eigenvectors} = \begin{bmatrix} 0.231 & -0.633 & 0.295 & 0.676 & 0.033 \\ -0.686 & 0.105 & 0.01 & 0.297 & 0.656 \\ -0.587 & -0.277 & -0.401 & 0.147 & -0.629 \\ -0.359 & -0.05 & 0.851 & -0.284 & -0.252 \\ -0.053 & -0.713 & -0.165 & -0.594 & 0.33 \end{bmatrix} \tag{6.8}$$

$$\text{Eigenvalues} = \begin{bmatrix} 1.462 & 1.333 & 1.013 & 0.696 & 0.596 \end{bmatrix} \tag{6.9}$$

An inspection of the eigenvalues shows that the first two principle components account for just over half of the variance in the data. In order to best represent the data in a 2 dimensional subspace we construct the 5x2 projection matrix W such that the columns of W are the eigenvectors that correspond to the two largest eigenvalues of the covariance matrix, so:

$$W = \begin{bmatrix} 0.231 & -0.633 \\ -0.686 & 0.105 \\ -0.587 & -0.277 \\ -0.359 & -0.05 \\ -0.053 & -0.713 \end{bmatrix} \tag{6.10}$$

We then project the points in X onto the subspace that is supported by it's two principle components, by calculating the product: $Y = X \times W$. Figure 6.9 shows a scatter plot of the two dimensional projection Y.

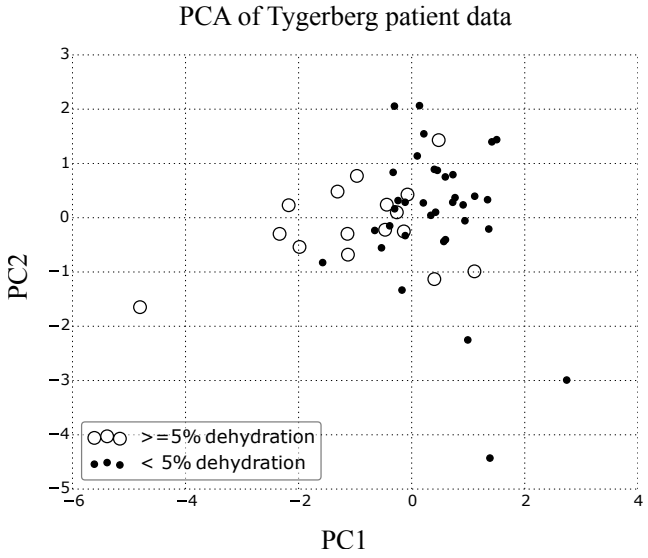

Figure 6.9: Scatter plot of the first two principle components of the Tygerberg patient dataset using the predictors in Equation 6.4

The data in the above plot shows the two components that represent the directions of maximum variance in the data and therefore the two dimensional representation that retains the most information of the original 5 dimensional dataset. However, if we are primarily interested in investigating the separability of the two classes (>= 5% dehydration and <5% dehydration) there is an alternative dimensionality reduction that might be more suited for the task. Linear discriminant analysis (LDA) is a method of projecting data onto a subspace, while maximising the separation between datasets with different labels. A very simple representation of the difference between the methods is given in Figure 6.10. Suppose you have two datasets (σ_1 and σ_2) in the two dimensional plane x, y, which you would like to reduce to one dimension. A PCA would ignore the class labels and project the data to the subspace ϕ in order to retain the maximum variance of the whole dataset σ^2_{1+2}. An LDA of the same data would try to maximize the distance between the means of the dataset while minimizing the variance of the sets and project the data to the subspace θ. It should be clear from Figure 6.10 that this analysis is worth pursuing for our separability exploration.

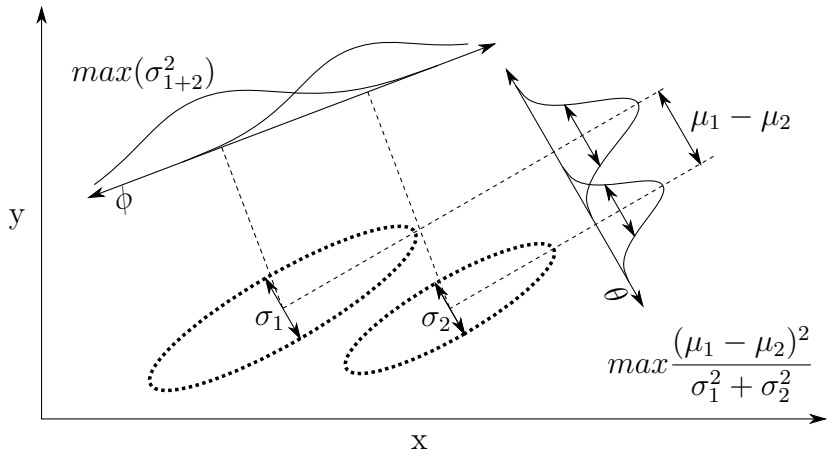

Figure 6.10: PCA and LDA projection

As with the PCA, the first step in the LDA analysis is to normalize the input vectors. The mean vectors are then calculated for each class.

$$m_i = \begin{pmatrix} \text{mean(digital capillary refill time)}_i \\ \text{mean(digital skin recoil time)}_i \\ \text{mean(maximum temperature gradient)}_i \\ \text{mean(age-normalised heart rate)}_i \\ \text{mean(age-normalised respiratory rate)}_i \end{pmatrix} \tag{6.11}$$

Where i = 1, 2 and represents the hydrated and dehydrated class respectively.

The within-class scatter matrix S_w is then calculated. S_w is the sum of the scatter matrix S_i for each class, where

$$S_i = \sum_{x \in D_i}^{N} (x - m_i)(x - m_i)^T \tag{6.12}$$

and m_i is the mean vector. The between class scatter matrix S_B is then calculated,

$$S_B = \sum_{i=1}^{c} N_i (m_i - m)(m_i - m)^T, \tag{6.13}$$

where m is the mean of all the samples and m_i and N_i are the mean and sample size of the respective classes.

The linear discriminants are then obtained by finding the eigenvectors and eigenvalues of the matrix $S_w^{-1} S_B$. The eigenvectors now represent the support for the data in the new feature subspace. To get the best representation of the data it is once again necessary to choose the eigenvectors with the largest eigenvalues. A quick look at the eigenvalues of the linear discriminants (eigenvalues = [0.874, 7.34e-17, 4.48e-17, 6.72e-18, 0.0]) shows that almost all of the interesting information is contained in the first linear discriminant. The transformation matrix W is then constructed again using the eigenvectors that

correspond to the largest two eigenvalues and the data is again projected to the new space Y. The projected data is shown in Figure 6.11

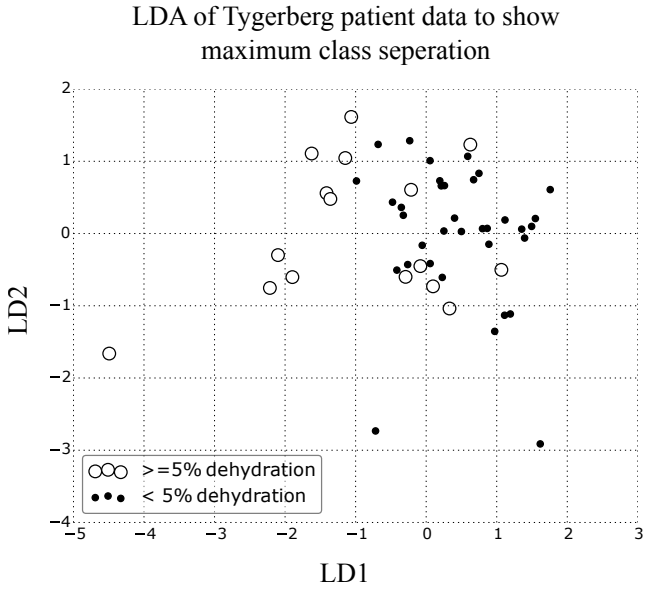

Figure 6.11: Scatter plot of the first two linear discriminants of the Tygerberg patient dataset using the predictors from Equation 6.4

The improvement in class separation for the LDA is not apparent from Figures 6.11 and 6.9. A side by side kernel density estimate of the first principle component and linear discriminant are shown in Figure 6.12 in an attempt to provide a more clear comparison of the two approaches.

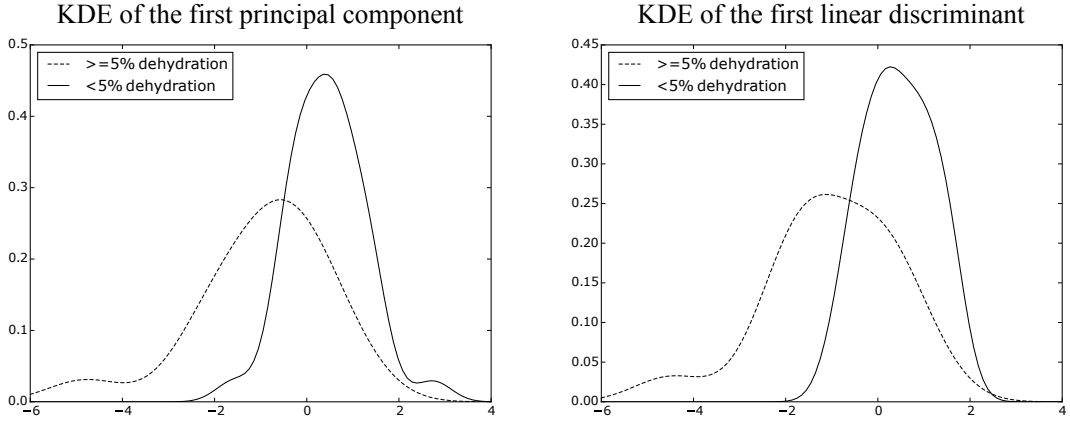

Figure 6.12: Kernel density estimate (KDE) to compare the class discriminatory power of the PCA and LDA

It is clear from Figure 6.12 that the gain in discriminatory power from using the LDA instead of the PCA is very small. The distributions in Figure

6.12 also provide us with an indication of what to expect from the final model. This result provides a fairly intuitive representation of the separability of the data and gives and indication of the performance that one could reasonably expect from the combined model. Before deriving the final model, there is one more analysis that would be instructive to illustrate the relative value of each predictor. An ordinary least squares linear regression of all the standardised predictors will show how much of the variance in the outcome is explained by each predictor. This along with the practical considerations associated with each predictor can then be used to assess which predictors are the most appropriate to include in the model.

The final model, that will be used to combine the predictors into one assessment parameter, will be derived by executing ordinary least squares multiple linear regression with the chosen parameters on the unstandardised data.

The result of the first linear regression will be a predicted hydration score, $\hat{y}$, (in arbitrary units) that is the weighted sum of the input predictors $x_1, x_2, ...x_n$ for each snapshot, as shown below.

$$\hat{y} = b_0 + b_1 x_1 + ... + b_n x_n \tag{6.14}$$

In the standardised form the regression equation becomes: $\hat{Y} = X.B$, where $\hat{Y}$ is the $N \times 1$ vector of the predicted scores for each snapshot (N = number of snapshots). X is the $N \times k$ matrix of predictors where the rows of X correspond to a sample or snapshot and the columns of X correspond to a predictor (k = number of predictors). B is the $k \times 1$ vector of model coefficients that relate the predictors to the output. One then calculates the coefficients in B such that the regression model minimizes the sum of the squared residuals $\sum(\hat{Y} - Y)^2$, where $\hat{Y}$ is the predicted hydration score for each snapshot and Y is the observed hydration score. To solve for B such that the residual is minimized the predicted hydration scores $\hat{Y}$, which is not yet available, is replaced with the observed hydration scores Y. To produce:

$$Y = X.B \tag{6.15}$$

B might not be square so both sides of the equation are pre multiplied by X^T so that the coefficient of B is invertible.

$$X^T Y = X^T X B \tag{6.16}$$

Both sides of this equation are then pre multiplied by the inverse of $X^T X$ to produce:

$$B = (X^T X)^{-1} X^T Y \tag{6.17}$$

A closer examination of the linear algebra shows that calculating B in this way produces coefficients that minimise the squared residual discussed above.

For the initial regression that is used to compare the values of each predictor all the available possible predictors were included. The predictor matrix therefore includes digital capillary refill time, digital skin recoil time, maximum temperature gradient, age-normalized heart rate, age-normalized respiratory rate, core temperature, eye appearance, fontanelle appearance, mucous membranes, neurological state and presence of tears.

The categorical variables from the clinical examinations were coded as 0,1 or 2 with increasing numbers corresponding to a condition that is associated with increased expected level of dehydration according to the literature (eg. eyes not sunken = 0, eyes sunken = 1). This means that any negative coefficients would indicate a correlation that contradicts the existing literature and is therefore probably not valid.

The results for the above analysis are summarised in Table 6.4

Table 6.4: Summary results for the multiple linear regression analysis on all the standardised input variables.

Marker	Coefficient	p-value	95%CI+	95%CI-
Device CRT	0.0070	0.955	-0.242	0.256
Device SRT	0.3699	0.014	0.080	0.660
Maximum temperature gradient	0.1229	0.384	-0.159	0.405
Normalized heart rate	0.2641	0.043	0.009	0.519
Normalized respiratory rate	0.0253	0.843	-0.231	0.282
Core temperature	-0.0278	0.836	-0.298	0.242
Sunken eyes	0.2708	0.073	-0.026	0.568
Dry mucous membrane	0.0845	0.588	-0.228	0.397
Sunken fontanelle	0.0352	0.789	-0.229	0.299
Tears present	-0.0959	0.553	-0.420	0.228
Neurologic status	0.2014	0.154	-0.079	0.481

From the above table one can see that the markers that correlated well with observed hydration level (in decreasing order) were digitally measured skin recoil time, eye appearance, age normalized heart rate, neurologic state, and maximum temperature gradient, of which only skin recoil time and normalized heart rate had a p-value of less then 0.05 which is the commonly accepted threshold for rejecting the null hypothesis. The R^2 value for the complete model is 0.478, which means that just under half of the variance in the output predictor can be explained by the model.

Of the hydration markers that seemed to explain some of the variance in the data only SRT and maximum temperature gradient were measured by the sensors that were developed during this study and can therefore form part of the final proposed sensor. The sensor fusion approach will therefore only include SRT and maximum temperature gradient as predictors for hydration

state. When deriving the sensor fusion algorithm the unstandardised values of the input and output variables will be used. This is to be able to use the formula with the raw data. If the model produced a perfect correlation, then the equation would relate the input variables in the units that they were rerecorded to the hydration state in percentage weight loss due to water loss. Since the correlation is fairly poor, it is better to interpret the output of the equation as a value that ranges roughly from 0 to 10 where higher values indicate a larger likely hood of dehydration. ROC curves can then be used to determine a good cut-off value for the predictor to determine a specific level of dehydration.

The results for the sensor fusion regression are tabulated below.

Table 6.5: Summary results for the multiple linear regression analysis on selected unstandardised markers.

Marker	Coefficient	p-value	95%CI+	95%CI-
Intercept	-1.5809	0.259	-4.362	1.200
Device SRT	60.6209	0.002	24.098	97.144
Maximum temperature gradient	6.7335	0.371	-8.270	21.737

The sensor fusion equation is therefore as follows:

$$\text{Hydration marker} = -1.58 + 60.6 \times \text{SRT} + 6.7 \times \text{Maximum temperature gradient} \tag{6.18}$$

The above equation has some interesting consequences. First of all notice that it has a non zero intercept since the input and output variables were not standardised. For the same reason the magnitude of the coefficients no longer represent how well the corresponding markers correlate with the output variable. Finally it is interesting to note the large coefficient of the skin recoil time. What this means is that for a one second increase in SRT the hydration marker would increase by around 60 for an average maximum temperature gradient. This would be a ludicrous result if one were to try and substitute digitally measured SRT with manually measured SRT (of which the resolution is often greater than 0.2 seconds) since the 'unit' for the output marker is still %weight loss due to water loss. This is however a realistic result for the population and markers of the study, since the skin recoil times that were recorded during the study using the developed device were very short and very consistent with a mean of 0.067 seconds and a standard deviation of 0.036 seconds.

The hydration marker that is determined by Equation 6.18 represents the best guess of hydration level from the available data. The marker is not reliable enough to make quantitative claims of hydration level. Fortunately however, to make treatment decisions the prediction does not have to be quantitative. The current hydration scales are used to divide patients into one of three or four hydration categories (eg. none, mild, moderate, severe), depending on the

treatment guidelines that are used. Therefore, it is best to treat the obtained result as the input to a binary classifier for a specific level of dehydration as was done with the previous hydration markers.

The ROC-curves that summarise the performance of the binary classifier for 5% and 10% dehydration, as well as the scatter plot of the hydration predictor vs. observed dehydration are shown in Figure 6.13. The AUC for the ROC curves below are 0.86 and 0.87 for 5% and 10% dehydration respectively.

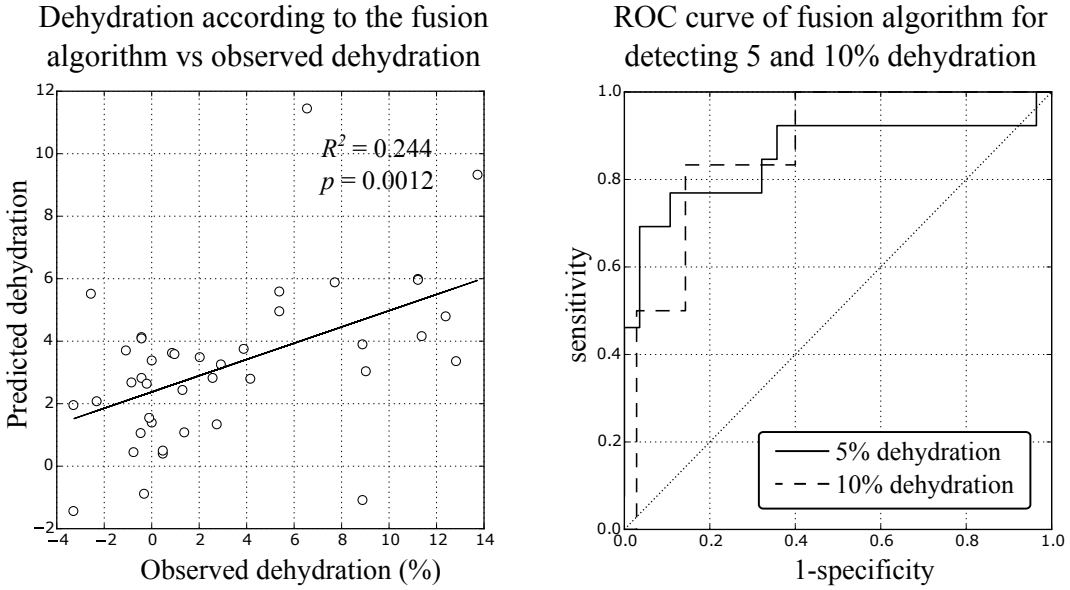

Figure 6.13: Scatter plot and ROC curves for regression based dehydration assessment

6.2.6 Notes on data conditioning for the sensor fusion analysis and clinical comparison

All the data presented above, that is related to a single sensor or marker was used verbatim without any outlier rejection or conditioning. The only data cleaning that was done, was performed during the marker extraction phase, where the samples with poor integrity were rejected.

For the results that comment on the performance of the clinical markers and clinical hydration scales, a separate data frame had to be constructed that consolidated the data from the clinical measurements with that of the normal assessment and digital measurement data. The datasets were combined using time and date stamps with an allowable time discrepancy of two hours. A lot can happen to a dehydrated patient in two hours, however a somewhat generous approach was required due to the sparsity of the clinical dataset.

The analysis that starts from Subsection 6.2.5 considered the combined performance of predictors. For this set the missing values were replaced by the means of the category values. A summary of the number of samples that were filled in are shown per analysis in Table 6.7 below, followed by a detailed breakdown of the same information.

Table 6.6: Summary of the filled in missing values, per analysis

Analysis	n	number of replaced samples	% Samples replaced
PCA/LDA analysis	255	57	22%
Standardised linear regression	561	204	36%
Unstandardised linear regression	204	43	21%

Table 6.7: Breakdown of the filled in missing values, per marker

Marker	number of replaced samples
Digital CRT	15
Digital SRT	5
Max temp gradient	6
Normalised heart rate	5
Normalised repertory rate	11
Core temperature	4
Sunken eyes	27
Mucous membranes	28
Sunken Fontanelle	29
Tears	31
Neurologic status	28

7. Discussion

7.1 Critical view of the CRT sensor

The CRT sensor was developed to allow for a more quantitative and objective measurement of CRT. From the literature survey it was clear that:

- CRT is a useful and proven clinical marker for the detection of dehydration.
- The assessment of CRT is not sufficiently standardised and there are not sufficient guidelines to adjust the interpretation of the CRT based on other parameters like temperature, blanching force and age which might affect the patient's CRT.
- The correct assessment of CRT is related to skill and experience of the physician and correct hydration assessments are less common in the settings where they are most needed.

The good: The sensor succeeded in providing good control over the application force and demonstrated that it is possible to control the balancing force with sufficient accuracy during the measurement by providing feedback to the operator. It also demonstrated a mechanism that allows the sensor to transition from pressure application to recording with minimal delay, by using the measured force to switch between modes and a spring loaded lever to move the pressure pad out of the way. The interface was simple and allowed for smooth operation during stressful and time sensitive situations, while allowing secure data handling. The sensor also successfully demonstrated that it is possible to extract a parameter that can be used to objectively describe blanching as shown in Figure 4.11, so that the refill time can be extracted by a simple and robust algorithm to ensure that the measurement of CRT is no longer subject to interpretation.

The bad: At 47 x 95 x 122 mm and weighing just under 300 grams the sensor was still far to large. Every attempt was made to keep the device as small as possible, however constraints related to the focal length and power requirements meant that the final device was still too bulky to allow for comfortable operation. The front end of the device was too wide and required the patient's shirts to be removed for the assessment. This was not a problem

for the infant study, since the measurements where made just before/after the patients were weighed (which required them to be completely undressed). It is however important for future devices since the effort required to dress and undress the patient will likely cause the sensor to be used incorrectly in an underserved clinical setting. The sensor was too sensitive to motion. Ill infants can be extremely restless. This restlessness translated to large artefacts in the patient data, which required a large portion of the CRT data to be discarded. The mouth of the sensor was not flexible enough. The mouth of the sensor offers about 10 mm of flex and the pressure pad protrudes 5 mm from the silicone cover which provides enough clearance between the rigid part of the mouth and the chest of the patient for most patient body types, however, this requires the sensor to be held perfectly horizontally. This was originally included by design to ensure that the sensor is correctly positioned during the recording phase. Unfortunately however, this required the device to be operated with two hands and lead the investigators to assume a somewhat stiff body position, this made the patients even more uncomfortable and greatly exacerbated the motion problem. The colour consistency of the recordings was not good enough. Upon analysing the CRT data it was found that a global colour change takes place in the scene, which was present despite the fact that (with great effort) all the automatic colour adjustments were turned off. These transients were also present in recordings of blank scenes which meant that they are likely the result of the lighting system or camera warming up at the start of the recording. The presence of these transients necessitated the use of an additional data conditioning step.

Lessons learned and suggestions for future versions

An actual camera is superfluous for this sensor. For this first attempt at quantifying blanching and capillary refill a camera was chosen to extract as much information as possible from the scene, since it was considered possible that features that are not obvious to the naked eye could be identified in the CRT footage. As it turns out the most robust representation of the level of blanching was the average lightness at the region of interest, which may as well have been recorded by a single, accurate, directional photo resistor or photo diode. Using a single value for the colour estimate at the region of interest will however require ideal lighting conditions and good lightness accuracy since there will be no background colour information to calculate a relative colour change as described in Section 4.1.4, unless a second light sensor is used. Motion will also need to be compensated for in a different manner, since no relative motion tracking information will be available. An ideal solution for the motion issue would be if the device could be made small and light enough so that it could move with the patient during operation. A smaller device that could be operated with one hand would also leave the other hand open to interact with and pacify the patient which would further reduce the prevalence of motion

artefacts. Finally the locking mechanism should be improved. For preliminary studies having a lock that needs to be actuated separately is probably still acceptable, however the current configuration requires the investigator to pull the locking lever and remove the downward force at the same time. If all the pressure is released to soon the camera will be triggered before the lever has had time to swing out of the way and if the downward force is removed after the lever is pulled the sensor will accelerate a few millimetres and bump against the patient. A lock that disengages after the pressure is removed will therefore be ideal for future versions of this sensor. It would also be desirable to redesign the swing arm in such a way that it is not mounted on the mouth of the sensor so that the entire light shield can be made from a more flexible material. It is also not necessary for the blanching site to be in the center of the scene. If for example it were moved right to the side of the sensor mouth it would be possible to record the CRT on the upper sternum between the collar of the patient's clothes and the patient's neck as demonstrated in Figure 7.1

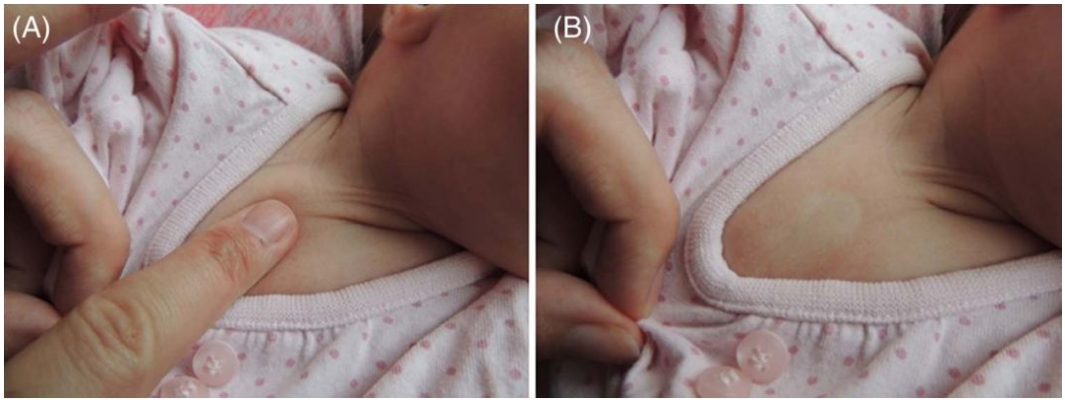

Figure 7.1: Demonstration of how to measure sternum CRT [17]

7.2 Critical view of the SRT sensor

The SRT sensor was developed to more accurately measure skin recoil time. The literature review shows that SRT is a valuable clinical dehydration marker. It is also apparent from the discussion in Section 3.1 and the results in Section 6.2.3 that it is not practical for physicians to manually determine SRT since, especially in infants, skin recoil happens so quickly.

The good: The sensor was able to capture skin recoil time with very high resolution. In fact, at the time that the study was conducted there was no evidence of any other study that tried to measure SRT this accurately. The sensor was fairly ergonomic and the combination of the directional lighting system and the 2 degree of freedom camera mount allowed for comfortable operation. As with the CRT sensor, the simple interface allowed for smooth

and hassle free operation in a stressful environment while facilitating secure data tagging and storage.

The bad: Manual time extraction. Determining the SRT from the video footage in a robust way is a fairly challenging computer vision problem and was not implemented in this study. Automatic SRT extraction will however be a requirement if SRT is ever to be assessed in practice in this way.

Lessons learned and suggestions for future versions

As discussed in the sensor development section, the purpose of this study was to investigate the basic feasibility of the first generation prototypes. For this reason not much time was dedicated to implementing the automated extraction of SRT. However, during the development of this sensor several methods of automatically extracting SRT from the raw footage were explored. They are discussed here for any future researchers that might want to pick up the gauntlet.

Using the generalised Hough Transform to measure skin fold

One of the first (and perhaps somewhat ambitious) designs for the SRT extraction involved drawing a line on the patient's skin and then conducting the SRT test in such a way that the skin fold is perpendicular to the line (This is why there is a line across the measurement site in Figure 4.14). The idea was that the image could then be sectioned using a canny (or similar) edge detector and that one could then try to extract the line using a generalised version of the Hough line transform, where the algorithm would 'vote' on the parameters of a parametric curve (like a cosine or inverse second or third order curve) in stead of a straight line. The parameters of the curve can then be tracked w.r.t time to determine when the fold has disappeared. The advantage of this method is that one also gets information about the shape of the fold, like fold thickness and fold length, which could contain other interesting information about the patient condition.

Using image contrast to detect recoil time

For the same exposure length images are more blurry when they move more quickly. It is therefore reasonable to expect that the contrast at the base of the skin fold will be high in the moments before the skin fold is released, low while the skin is recoiling and high when the skin has recoiled. This information might then be used to automatically extract skin recoil time.

Dense optical flow

Arguably the most obvious way of extracting SRT from the raw footage is to calculate the optical flow vectors in the scene and determine when the motion at the skin fold starts and stops. Finding the optical flow in a scene is computationally very expensive, but in some cases the optical flow vectors are already available, and it is useful to know to look for them. For example, the Raspberry Pi, on which this sensor is based, is itself based on a BCM2835 SoC, which can encode .h264 video as it is recorded. Since the encoding is done on the GPU, it is possible to access parts of the encoding pipeline while recording. One of the parts of the recording pipeline that one can access on the Raspberry Pi is the motion estimation block. The motion estimation block is used for the video encoding and contains the motion vectors of the mark-blocks (16x16 pixel blocks that make up the scene). This motion vector calculation is implemented in hardware and is therefore calculated in an extremely quick and efficient manner.

7.3 Critical look at the thermal profile sensor

Core to peripheral temperature gap is often cited in literature as a useful parameter for the assessment of hydration status. It is not as powerful, according to most studies, as either CRT or SRT, but it still has established diagnostic value. The purpose of the thermal profile sensor that was developed during this study was to determine if there is more information that can be extracted from the skin temperature profile of infants that can be related to their hydration level.

The good: The FLIR E-60 camera that was used during the study captured all the available skin temperature information with good resolution.

The bad: The algorithm that extracted the hydration marker from the temperature profile only extracted and evaluated the maximum temperature gradient along a provided path. It would have been interesting to see if one could extract other parameters from the thermal images like 'blotchiness' or temperature change w.r.t. time after a disturbance which one could relate to peripheral perfusion or hydration level.

Lessons learned and suggestions for future versions:

The use of thermal imaging in medicine is still in its' infancy. This is partly due to the fact that until recently long-wave thermal imaging sensors were prohibitively expensive. This seems to be rapidly changing. Seek thermal has recently introduced a 206x156 pixel, long wave infrared, Vanadium Oxide Microbolometer based thermal image sensor which it sells for $250-$300. It

therefore seems likely that many new ways will be discovered to use this technology for patient assessment. While there was a lot of collaboration between BERG and Tygerberg hospital during this project, from the point of inception up to the validation study, it is the opinion of the investigator that even more input from medical experts would be valuable. Especially since the thermal profile sensor was built to study a medical/physiological hypothesis that was proposed by an engineering student.

7.4 Limitations of the Tygerberg Study

The primary limitation of the Tygerberg study is the sample size and sparsity of the clinical assessments. The study was conducted over a period of four non-consecutive weeks between November 2014 and February 2015. The initial plan was to complete the study in November, since according to the Western Cape department of Health, this was the start of dehydration season. After the first two weeks however, only five patients were enrolled in the study. It was decided that this was not cost effective, because study design required the investigators to be in the ward 24 hours a day so that the first hydration assessments can be made within the first hour of admission to the ward. The study was therefore paused and resumed in February which is the peak of the Western Cape dehydration season. After two more weeks of testing a further five patients were enrolled in the study. The low enrolment numbers, combined with the large prevalence of motion artefacts in the CRT sensor data meant that it made more sense to end the study sooner. At the end of February it was decided to conclude the study and use the available data to make an assessment of the sensor feasibility as well as develop recommendations for future projects. While the digital marker assessments and nurse assessments are fairly complete, the clinical assessments are woefully sparse. This is partly due to the fact that the ward was constantly understaffed in the period that the assessments were made.

7.5 Critical look at the performance of the sensors

CRT sensor

The performance of the CRT sensor in the infant trials was marred by factors of the final setting that were not sufficiently addressed in the original design. No interesting conclusions can therefore be made about the clinical value of a better measurement of CRT in a clinically relevant setting, because no such data was collected.

SRT sensor

The skin recoil time sensor provided some really interesting results. The first thing that one can see from Figure 6.6 is that the average recoil time as determined by the camera and operator combination is very short (mean of 0.08 seconds). It could be that using this method of SRT assessment introduces a systematic bias by only measuring the first part of the skin recoil event. Whatever the case may be, the data seems to suggest that there is a fair correlation between the measured SRT and the hydration state. While the correlation in Figure 6.6 might not seem overwhelmingly strong, the ROCs obtained from the data for predicting 5% and 10% dehydration are very promising. In fact a sensitivity and specificity of greater than 80% for a single non invasive measurement is frankly unheard of in the field. While it would be unrealistic to expect that these results could be exactly replicated, the evidence makes a good case for the continued exploration of high resolution SRT assessments for the assessment of hydration level in infants.

Thermal profile sensor

While the ROC curve in Figure 6.8 seems to suggest fair performance, the thermal profile sensor validation data set suffers from the same ailments as that of the others, that is, small sample size. With a sample of only 10 patients and 51 point estimates one would really need exceptional performance to peak the interest of the greater scientific community who are (understandably) likely to disregard this particular combination of performance and statistical backing as "not particularly interesting". The fact that this method is based on a fairly intuitively derived hypothesis might (hopefully) inspire future researchers to investigate the method in greater detail, by investigating other marker extraction algorithms and grounding their results in a larger study.

Sensor fusion approach

The sensor fusion method, showed that it is possible to improve the prediction from the sensors if one is able to find complementary markers, i.e. markers that correlate well to the output, but not to one another. It is however important to note that as a result of the small data set, the same data was used to derive and assess the performance of the fusion approach, which means that the sensor fusion model will likely perform worse on a different data set that was captured in the same way. Because of the small sample size the fusion approach was limited to OLS multiple linear regression. It is however likely that there are many non-linear interactions between the various physiological markers that are not captured by the proposed model and it would be very interesting to see what relationships could be identified in future projects.

A critical look at the proposed final product (system)

The proposed sensor system does not make use of any particularly novel technology and is unlikely to revolutionise the way in which hydration is measured, it does however outline a system the can feasibly improve the way that hydration estimations are made in a underserved setting by introducing a more ridged test methodology and by standardising many of the factors that confound the measurement of hydration markers. The success of the proposed system will therefore likely depend on the price and ergonomics of the final version, should it get that far. While the presented results may seem inconclusive it is important to remember that the study tested first generation sensors in a difficult setting and that the purpose of the tests was to prove the basic feasibility of the sensors, which it did. Evidence of the clinical value of the markers is abundant in the literature and the only thing that the clinical study had to prove was that it is possible to get repeatable results in a clinical setting. While the results are far from conclusive, this seems to be the case according to the noise ratios and test-retest performance of the sensors as discussed in Section 6.1.

8. Conclusion

Three sensors for the non-invasive assessment of infant dehydration were proposed. Two of which are based on accepted clinical markers. The first two sensors were then designed, built and tested. The performance of the third sensor (Thermal Profile Sensor) was emulated using a thermal camera and custom post processing scripts. The devices were tested in two separate clinical studies involving adults and infants. The purpose of the adult study was to prove that the devices were capable of making repeatable measurements in a research setting. This was proven to be the case for the adult study.

The infant study was conducted to test whether the sensors can be used in a stressful clinical setting and to determine whether a correlation can be found between the measured hydration markers and hydration state. The CRT sensor did not fare well in the infant study. While it was ergonomic enough to use, many of the samples had to be rejected because of poor data integrity and no real correlation could be found between the hydration state and CRT. This, however, is likely due to the small sample size and poor data integrity since there is overwhelming proof in literature that CRT is indeed a good indication of hydration. In contrast to the CRT sensor, the SRT sensor faired quite well. It's size did not seem to upset the infants as much as that of the CRT sensor and the measured SRT correlated well to the hydration states of the infants. The thermal profile sensor performed better than the CRT sensor, but not as well as the SRT sensor, and as with the other sensors, no firm conclusions can be drawn about it's clinical performance due to the small sample size of the infant study.

Apart from the conclusions that can explicitly be drawn from the data that is presented in the preceding sections, many important lessons that are related to the project goals were learned while conducting the study. Arguably the most important lesson is that dehydrated infants are by no means stationary, and how the patients will perceive and interact with the devices should be one of the primary design considerations for future iterations.

The results obtained in the two studies are a mixed bag, however, given that infant mortality due to diarrhoea related dehydration in underserved settings is such a massive burden, all avenues of addressing the problem should be explored. The proposed sensors contain no expensive components and if a version of these sensors can be designed, that perform as well as the literature suggests,

it should be possible to bring this solution to the poorest and most affected populations. The scalability of the solution, combined with the severity of the problem, the literature's support for the markers and the results presented here, should therefore be enough to inspire further investigation into the approach.

List of References

[1] Friedman, J.N., Goldman, R.D., Srivastava, R. and Parkin, P.C.: Development of a clinical dehydration scale for use in children between 1 and 36 months of age. *The Journal of pediatrics*, vol. 145, no. 2, pp. 201–7, August 2004. ISSN 0022-3476.
Available at: `http://www.ncbi.nlm.nih.gov/pubmed/15289767`

[2] Boschi-Pinto, C., Velebit, L. and Shibuya, K.: Estimating child mortality due to diarrhoea in developing countries. *Bulletin of the World Health Organization*, vol. 86, no. 9, pp. 710–717, 2008. ISSN 00429686.

[3] *World Health Statistics 2015*. WHO Press, Geneva, 2015. ISBN 978 92 4 156488 5.
Available at: `http://apps.who.int/iris/bitstream/10665/170250/1/9789240694439_eng.pdf?ua=1&ua=1`

[4] Brook, C.G.: Determination of body composition of children from skinfold measurements. *Archives of disease in childhood*, vol. 46, no. 246, pp. 182–4, April 1971. ISSN 1468-2044.
Available at: `http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1647464&tool=pmcentrez&rendertype=abstract`

[5] Finberg, L.: Dehydration in infancy and childhood. *Pediatrics in review / American Academy of Pediatrics*, vol. 23, no. 8, pp. 277–82, August 2002. ISSN 1526-3347.
Available at: `http://www.ncbi.nlm.nih.gov/pubmed/12154234`

[6] Huang, L.H. and Corden, T.E.: Medscape, Dehydration Refrince. *Dehydration - Medscape Reference*, 2012. Last accesed on 2013-06-06.
Available at: `http://emedicine.medscape.com/article/906999-overview#showall`

[7] Cheuvront, S.N., Ely, B.R., Kenefick, R.W. and Sawka, M.N.: Biological variation and diagnostic accuracy of dehydration assessment markers. *American Journal of Clinical Nutrition*, vol. 92, no. 3, pp. 565–573, 2010. ISSN 00029165.

[8] Armstrong, L.E.: Assessing hydration status: the elusive gold standard. *Journal of the American College of Nutrition*, vol. 26, no. 5 Suppl, pp. 575S–584S, October 2007. ISSN 0731-5724.
Available at: http://www.ncbi.nlm.nih.gov/pubmed/17921468

[9] Armstrong, L.E.: Hydration Assessment Techniques. *Nutrition Reviews*, vol. 63, no. 6, 2008.

[10] Cheuvront, S.N., Ely, B.R., Kenefick, R.W. and Sawka, M.N.: Biological variation and diagnostic accuracy of dehydration assessment markers. *The American journal of clinical nutrition*, vol. 92, no. 3, pp. 565–573, 2010.

[11] Bourque, C.W.: Central mechanisms of osmosensation and systemic osmoregulation. *Nature reviews. Neuroscience*, vol. 9, no. 7, pp. 519–31, July 2008. ISSN 1471-0048.
Available at: http://www.ncbi.nlm.nih.gov/pubmed/18509340

[12] Steiner, M.J. and Dewalt, D.A.: Is This Child Dehydrated ? *The Rational Clinical Examination*, vol. 291, no. 22, pp. 2746–2754, 2004.

[13] Gorelick, M.H., Shaw, K.N. and Murphy, K.O.: Validity and reliability of clinical signs in the diagnosis of dehydration in children. *Pediatrics*, vol. 99, no. 5, p. E6, 1997. ISSN 0031-4005.

[14] Pringle, K., Shah, S.P., Umulisa, I., Mark Munyaneza, R.B., Dushimiyimana, J., Stegmann, K., Musavuli, J., Ngabitsinze, P., Stulac, S. and Levine, A.C.: Comparing the accuracy of the three popular clinical dehydration scales in children with diarrhea. *International Journal of Emergency Medicine*, vol. 4, no. 1, p. 58, 2011. ISSN 1865-1380.
Available at: http://www.intjem.com/content/4/1/58

[15] Falszewska, A., Dziechciarz, P. and Szajewska, H.: The Diagnostic Accuracy of Clinical Dehydration Scale in Identifying Dehydration in Children With Acute Gastroenteritis: A Systematic Review. *Clinical Pediatrics*, vol. 53, pp. 1181–1188, 2014. ISSN 0009-9228.
Available at: http://cpj.sagepub.com/cgi/doi/10.1177/0009922814538493

[16] Lima, A.P., Beelen, P. and Bakker, J.: Use of a peripheral perfusion index derived from the pulse oximetry signal as a noninvasive indicator of perfusion. *Critical care medicine*, vol. 30, pp. 1210–1213, 2002. ISSN 0090-3493.

[17] King, D., Morton, R. and Bevan, C.: How to use capillary refill time. *Archives of disease in childhood. Education and practice edition*, pp. 1–6, November 2013. ISSN 1743-0593.
Available at: http://www.ncbi.nlm.nih.gov/pubmed/24227793

[18] Pickard, A., Karlen, W. and Ansermino, J.M.: Capillary refill time: Is it still a useful clinical sign? *Anesthesia and Analgesia*, vol. 113, no. 1, pp. 120–123, 2011. ISSN 00032999.

[19] Pandey, A. and John, B.M.: Capillary refill time. is it time to fill the gaps? *Medical Journal Armed Forces India*, vol. 69, pp. 97–98, 2013. ISSN 03771237.

[20] Bouguet, J.-Y.: Pyramidal Implementation of the Lucas Kanade Feature Tracker Description of the algorithm. *In Practice*, vol. 1, no. 2, pp. 1–9, 1999. ISSN 0966842X. 3629719.
Available at: `http://trac.assembla.com/dilz_mgr/export/272/doc/ktl-tracking/algo_tracking.pdf`

[21] Jauregui, J., Nelson, D., Choo, E., Stearns, B., Levine, A.C., Liebmann, O. and Shah, S.P.: External validation and comparison of three pediatric clinical dehydration scales. *PLoS ONE*, vol. 9, no. 5, pp. 1–6, 2014. ISSN 19326203.

[22] Illner, K., Brinkmann, G., Heller, M., Bosy-Westphal, a. and Müller, M.J.: Metabolically active components of fat free mass and resting energy expenditure in nonobese adults. *American journal of physiology. Endocrinology and metabolism*, vol. 278, no. 2, pp. E308–E315, 2000. ISSN 0193-1849.

[23] Gallagher, D., Belmonte, D., Deurenberg, P., Wang, Z., Krasnow, N., Pi-Sunyer, F.X. and Heymsfield, S.B.: Organ-tissue mass measurement allows modeling of REE and metabolically active tissue mass. *The American journal of physiology*, vol. 275, no. 2 Pt 1, pp. E249–E258, 1998. ISSN 0002-9513.

[24] Duggan, C., Refat, M., Hashem, M., Wolff, M., Fayad, I. and Santosham, M.: Journal of Pediatric Gastroenterology & Nutrition : How Valid Are Clinical Signs of Dehydration in Infants ? How Valid Are Clinical Signs of Dehydration in Infants ? : Journal of Pe ... *Journal of Pediatric Gastroenterology & Nutrition*, vol. 22, no. 1, pp. 56–61, 1996.

Appendices

A. CAD Drawings

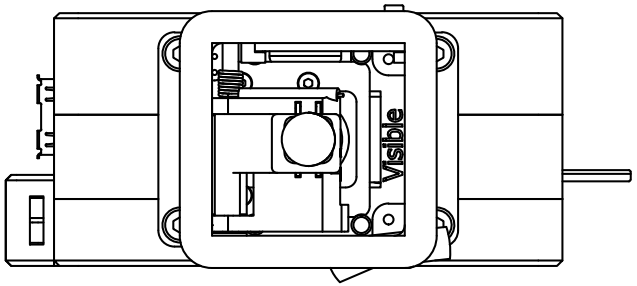

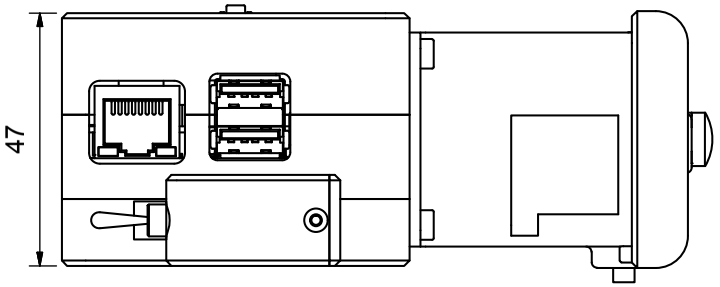

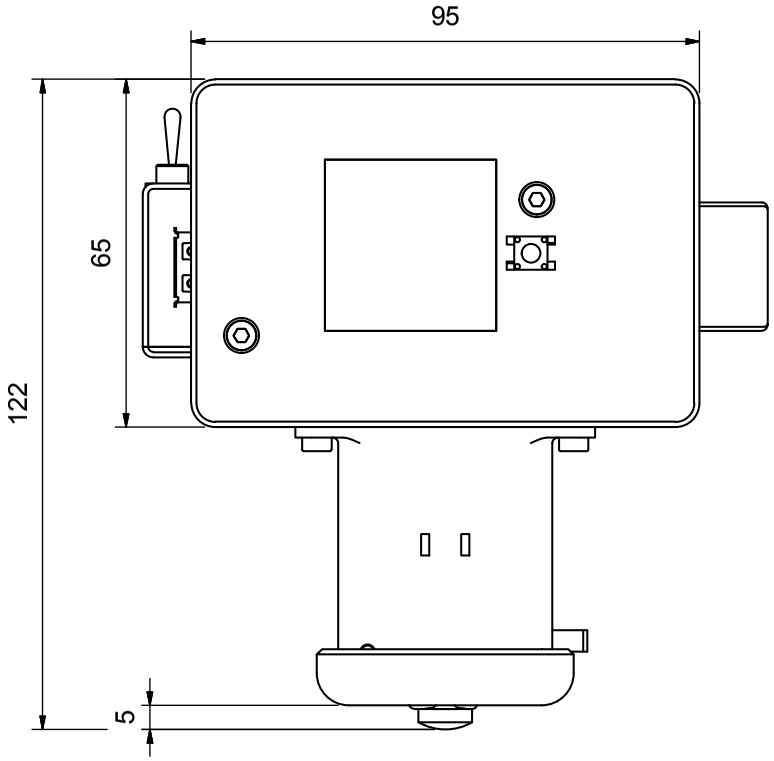

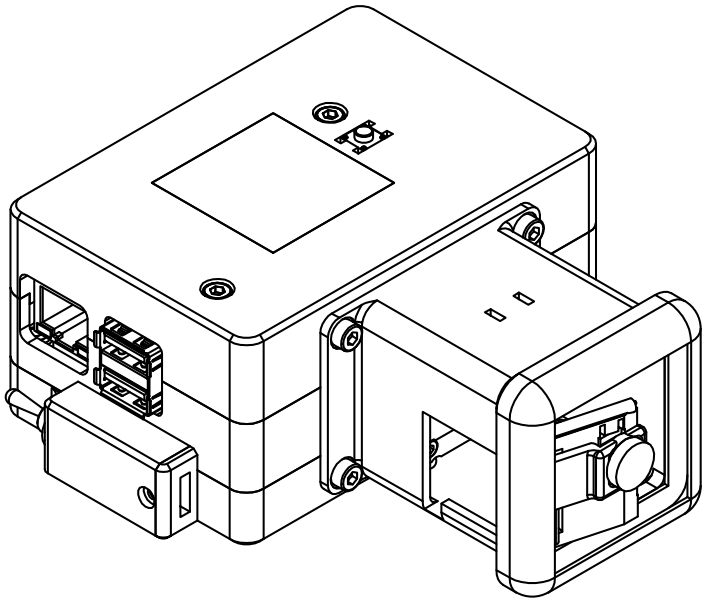

ITEM	BESKRYWING	AANTAL	MATERIAAL / SPESIFIKASIES

UNIVERSITEIT VAN STELLENBOSCH

SKAAL OP A 3 1:1		TITEL: CRT sensor	
MATE IN mm			
STUDENTE No. 15629635	TEKENAAR Eduard Kieser	NAGESIEN	DATUM 23-7-2015
VEL No. VAN VELLE			No.

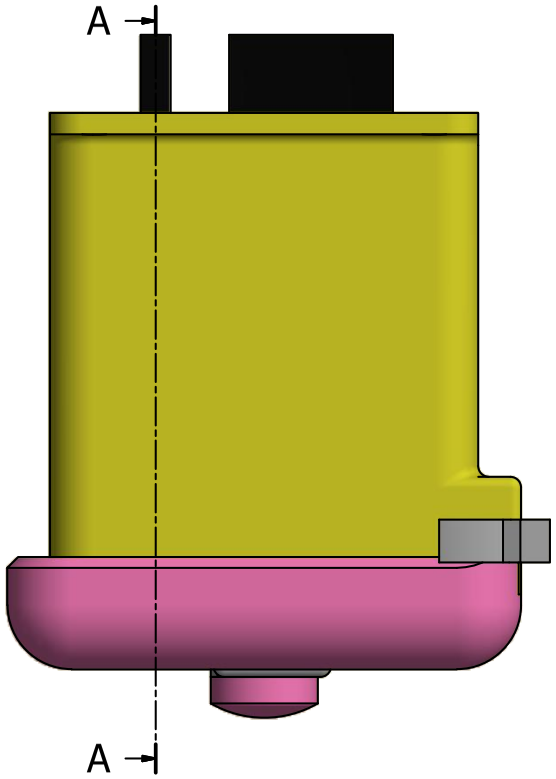

A-A (2)

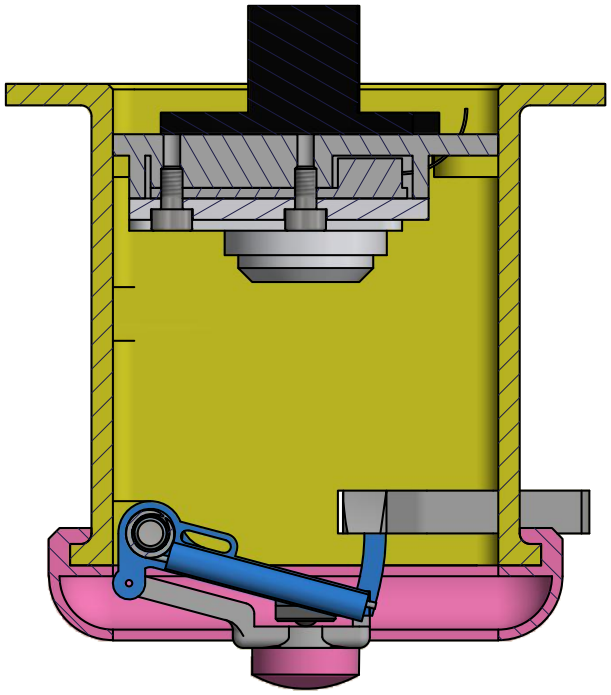

B-B (2)

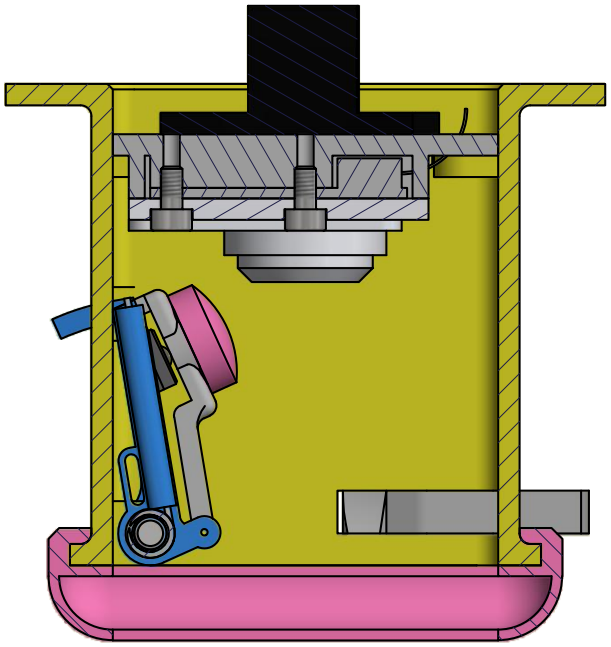

ITEM	BESKRYWING	AANTAL	MATERIAAL / SPESIFIKASIES

UNIVERSITEIT VAN STELLENBOSCH

SKAAL OP A 3 2:1			TITEL:	Section of light shield to show swing arm and force sensor	
MATE IN mm					
STUDENTE No. 15629635	TEKENAAR Eduard Kieser	NAGESIEN	DATUM 23-7-2015	VEL No. VAN VELLE	No.

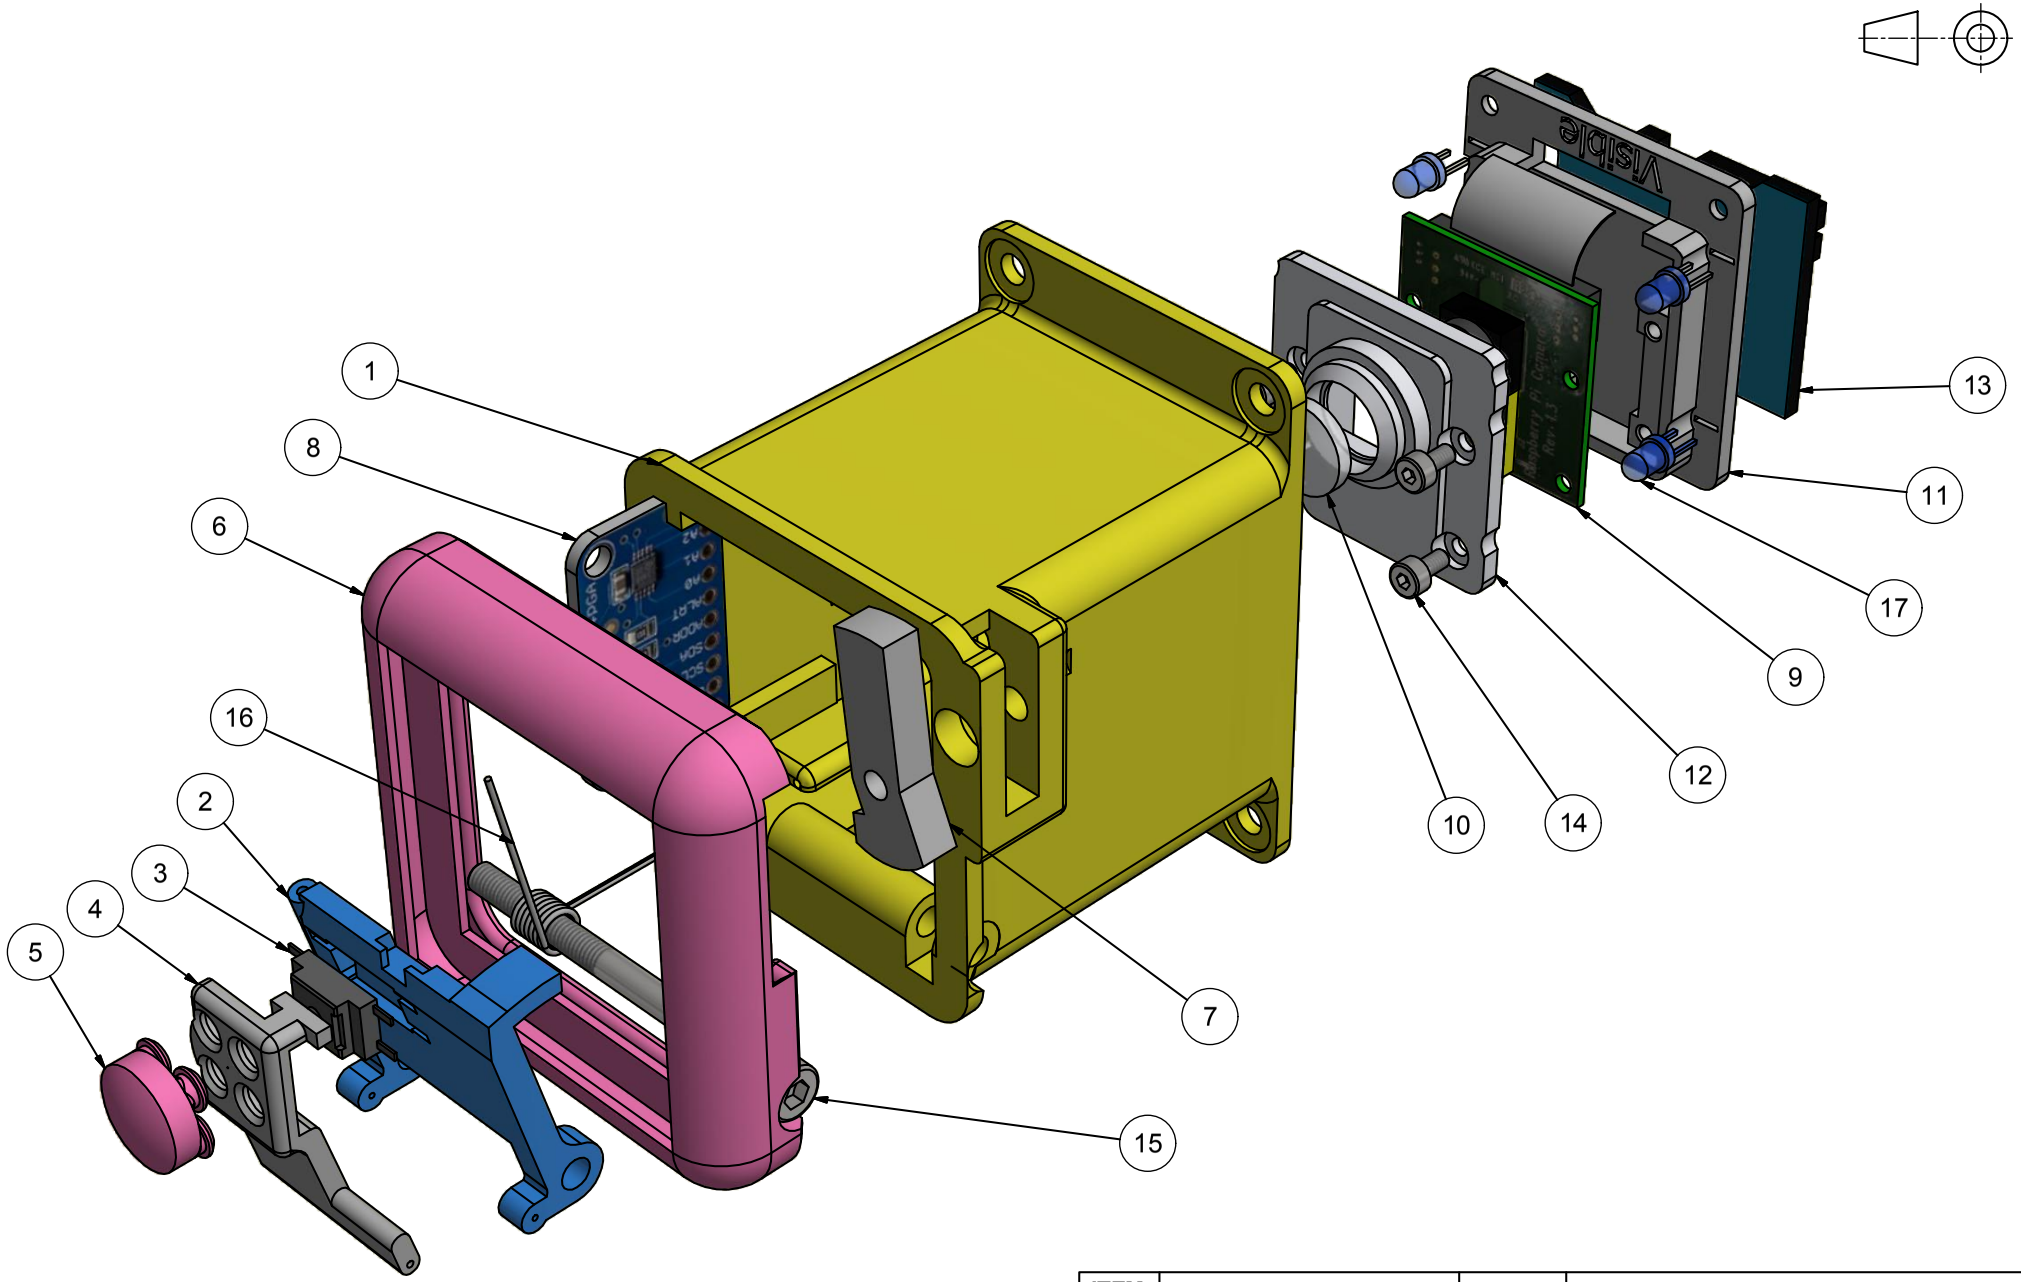

ITEM	BESKRYWING	AANTAL	MATERIAAL / SPESIFIKASIES

UNIVERSITEIT VAN STELLENBOSCH

SKAAL OP A 3		TITEL:	CRT light shield exploded view
MATE IN			
STUDENTE No. 15629635	TEKENAAR Eduard Kieser	NAGESIEN	DATUM 23-7-2015
VEL No.	VAN	VELLE	No.

ITEM	BESKRYWING	AANTAL	MATERIAAL / SPESIFIKASIES
17	60 degree White LEDs	4	
16	Tortion Spring	1	Spring Steel
15	Hexagon Socket Head Cap Screw	1	Stainless Steel
14	Hexagon Socket Head Cap Screw	4	Stainless Steel
13	Lighting Board	1	
12	Camera Cover and Lens Mount	1	3D Printed Acrylic
11	Camera Base Plate	1	3D Printed Acrylic
10	Lens	1	Coated N-BK7
9	Camera Module	1	
8	Analoug to Digital Converter	1	
7	Locking Lever	1	3D Printed Acrylic
6	Silicone Contact Surface	1	Silicone
5	Pressure Pad	1	Silicone
4	Pressure Pad Mount	1	3D Printed Acrylic
3	Force Sensor	1	
2	Swing Arm	1	3D Printed Acrylic
1	Light Shield	1	3D Printed Acrylic

UNIVERSITEIT VAN STELLENBOSCH

SKAAL OP A	TITEL: CRT light shield parts list	
MATE IN		
STUDENTE No. / TEKENAAR / NAGESIEN	DATUM	VEL No. VAN VELLE No.

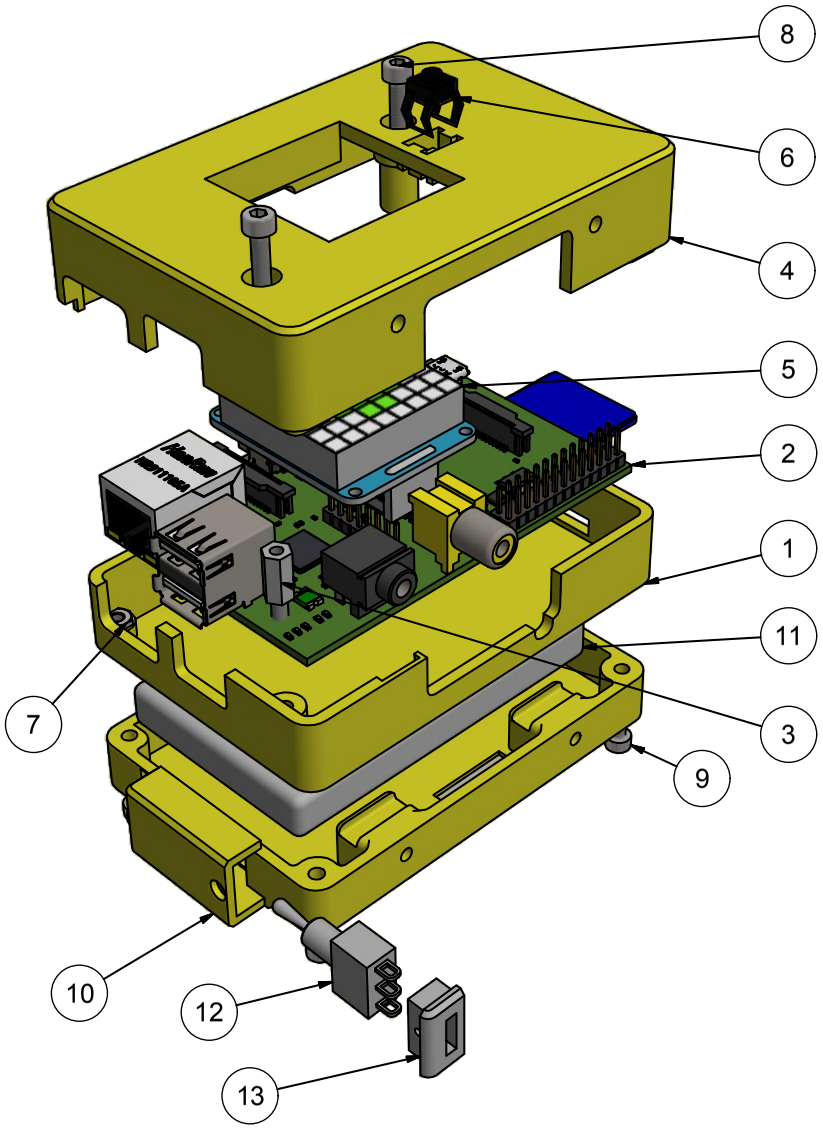

ITEM	BESKRYWING	AANTAL	MATERIAAL / SPESIFIKASIES
13	Charging Plug Bracket	1	3D Printed Acrylic
12	Power Switch	1	
11	Battery	1	
10	Battery Bay	1	3D Printed Acrylic
9	Hexagon Socket Head Cap Screw	4	Stainless Steel
8	Hexagon Socket Head Cap Screw	2	Stainless Steel
7	Hexagon Head Bolts	4	Stainless Steel
6	Push Button	1	
5	Display	1	
4	Houding Cover	1	3D Printed Acrylic
3	25mm stand off	2	Nylon
2	Raspberry Pi Model B	1	
1	Housing Base	1	3D Printed Acrylic

UNIVERSITEIT VAN STELLENBOSCH

SKAAL OP A 3		TITEL:	Controller Housing for CRT/SRT sensor
MATE IN			
STUDENTE No. 15629635	TEKENAAR Eduard Kieser	NAGESIEN	
DATUM 23-7-2015		VEL No. VAN VELLE	No.

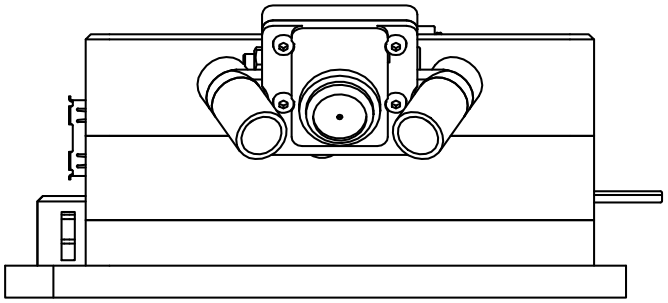

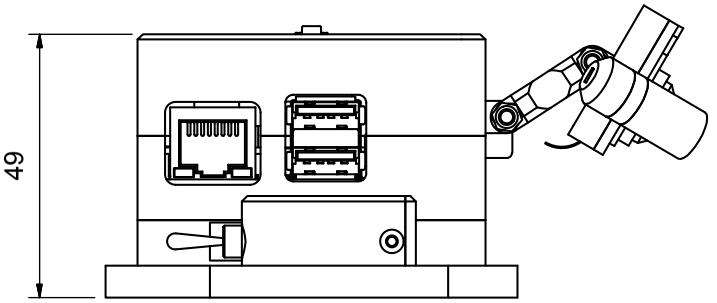

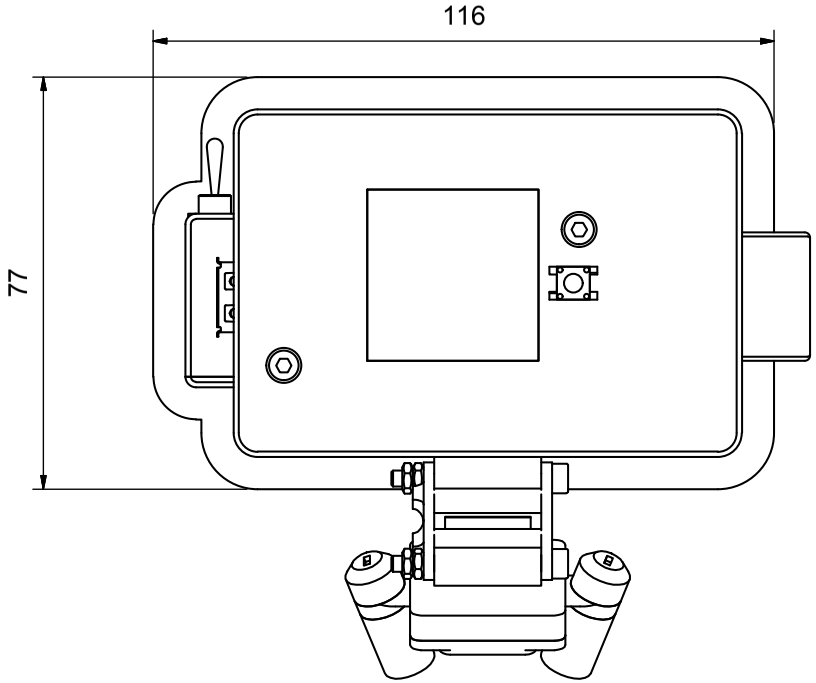

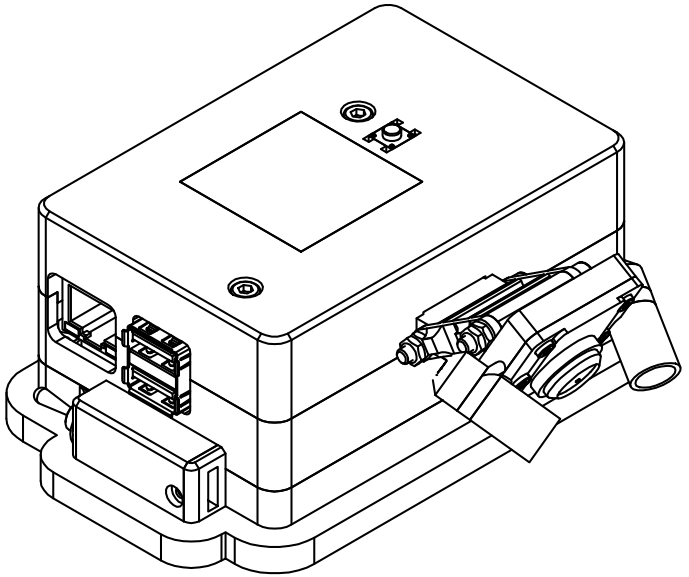

ITEM	BESKRYWING	AANTAL	MATERIAAL / SPESIFIKASIES

UNIVERSITEIT VAN STELLENBOSCH		SKAAL OP A 3 1:1 MATE IN mm	TITEL: SRT sensor
STUDENTE No. 15629635	TEKENAAR Eduard Kieser	NAGESIEN	DATUM 23-7-2015

VEL No. VAN VELLE No.

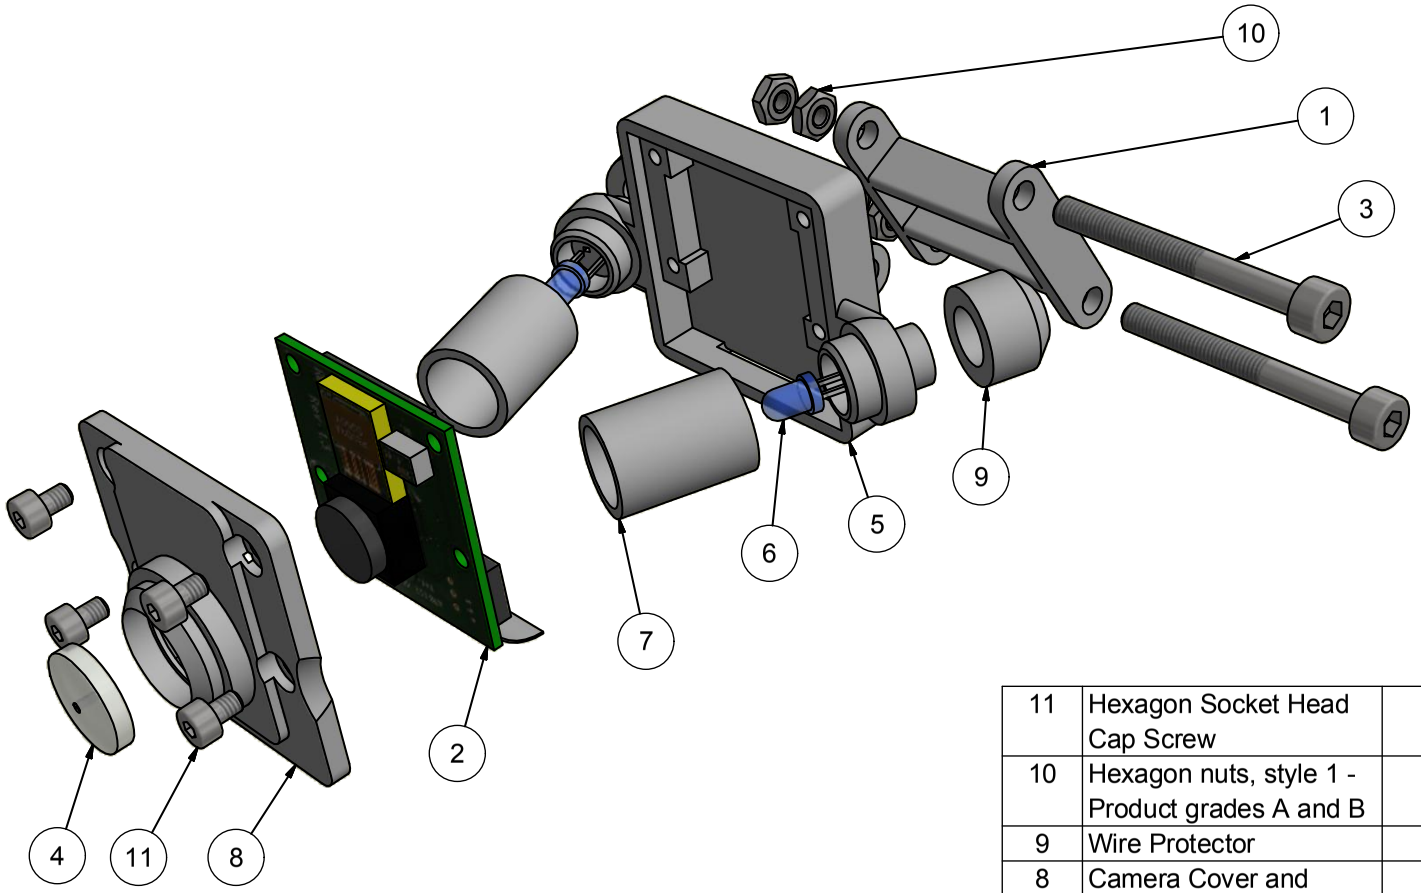

ITEM	BESKRYWING	AANTAL	MATERIAAL / SPESIFIKASIES
11	Hexagon Socket Head Cap Screw	4	Stainless Steel
10	Hexagon nuts, style 1 - Product grades A and B	4	Stainless Steel
9	Wire Protector	2	3D Printed Acrylic
8	Camera Cover and Lens Mount	1	3D Printed Acrylic
7	Light Funnel	2	3D Printed Acrylic
6	60 degree White LEDs	2	
5	Camera Backplate and Lighting Mount	1	3D Printed Acrylic
4	Lens	1	Coated N-BK7
3	Hexagon Socket Head Cap Screw	2	Stainless Steel
2	Camera Module	1	
1	2 DOF Camera Arm	1	3D Printed Acrylic

UNIVERSITEIT VAN STELLENBOSCH

STUDENTE No. 15629635	TEKENAAR Eduard Kieser	NAGESIEN	SKAAL OP A 3 MATE IN	TITEL: Exploded view of SRT camera	
			DATUM 23-7-2015	VEL No. VAN VELLE	No.

B. Schematics

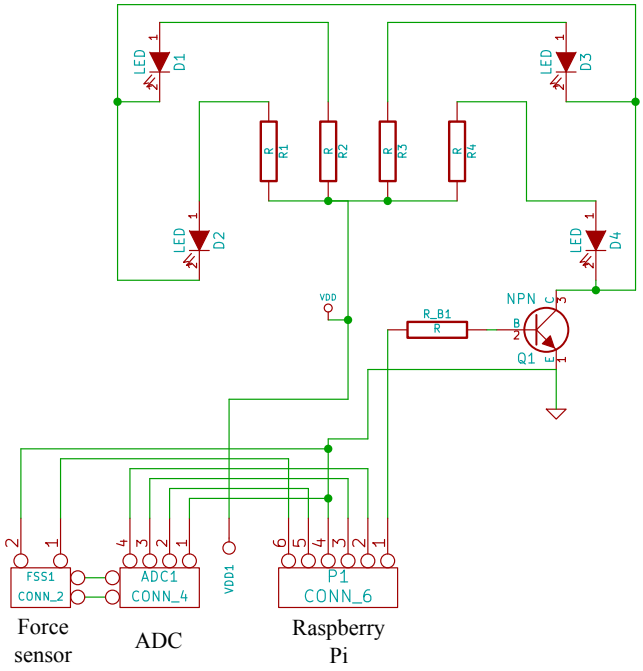

Figure B.1: Schematics of the lighting board, which includes a bridge harness to easily connect the force sensor and ADC to the controller.

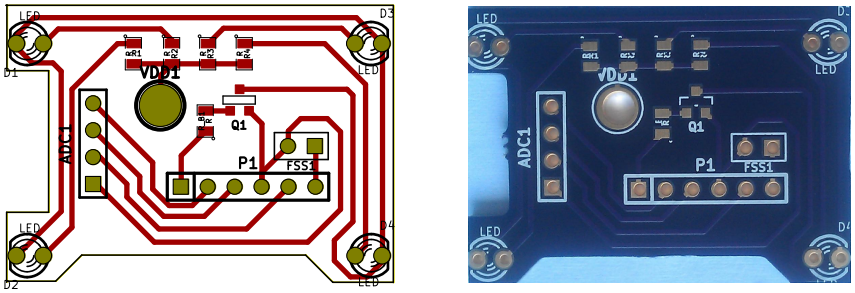

Figure B.2: Track layout (a) and image (b) of the lighting board

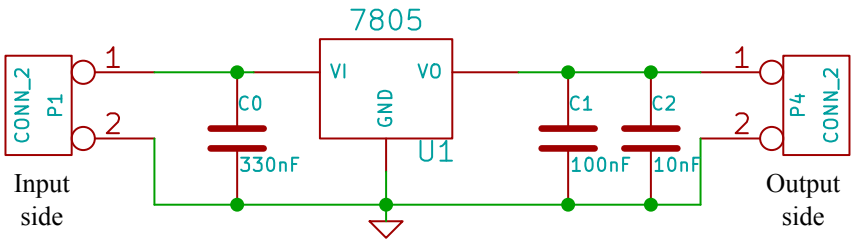

Figure B.3: Schematics of the 5V power regulator that was used in both the CRT and SRT sensor.

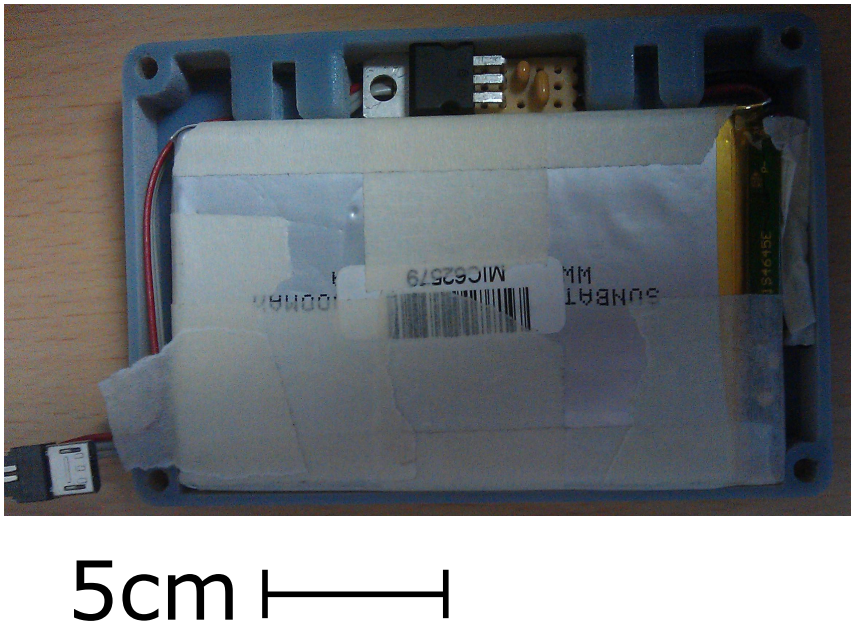

Figure B.4: The 7.4V, 2000MAh battery and regulator that was used in the CRT and SRT sensor.

C. Data Extraction Software Interface Screen-Shots

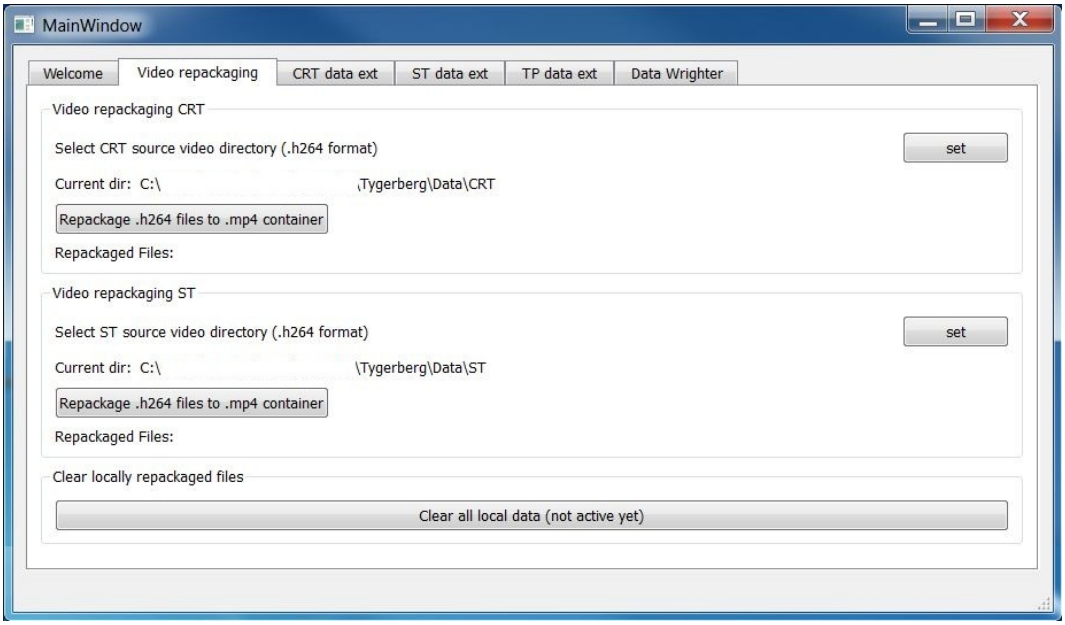

Figure C.1: Data extraction application, video repackaging interface

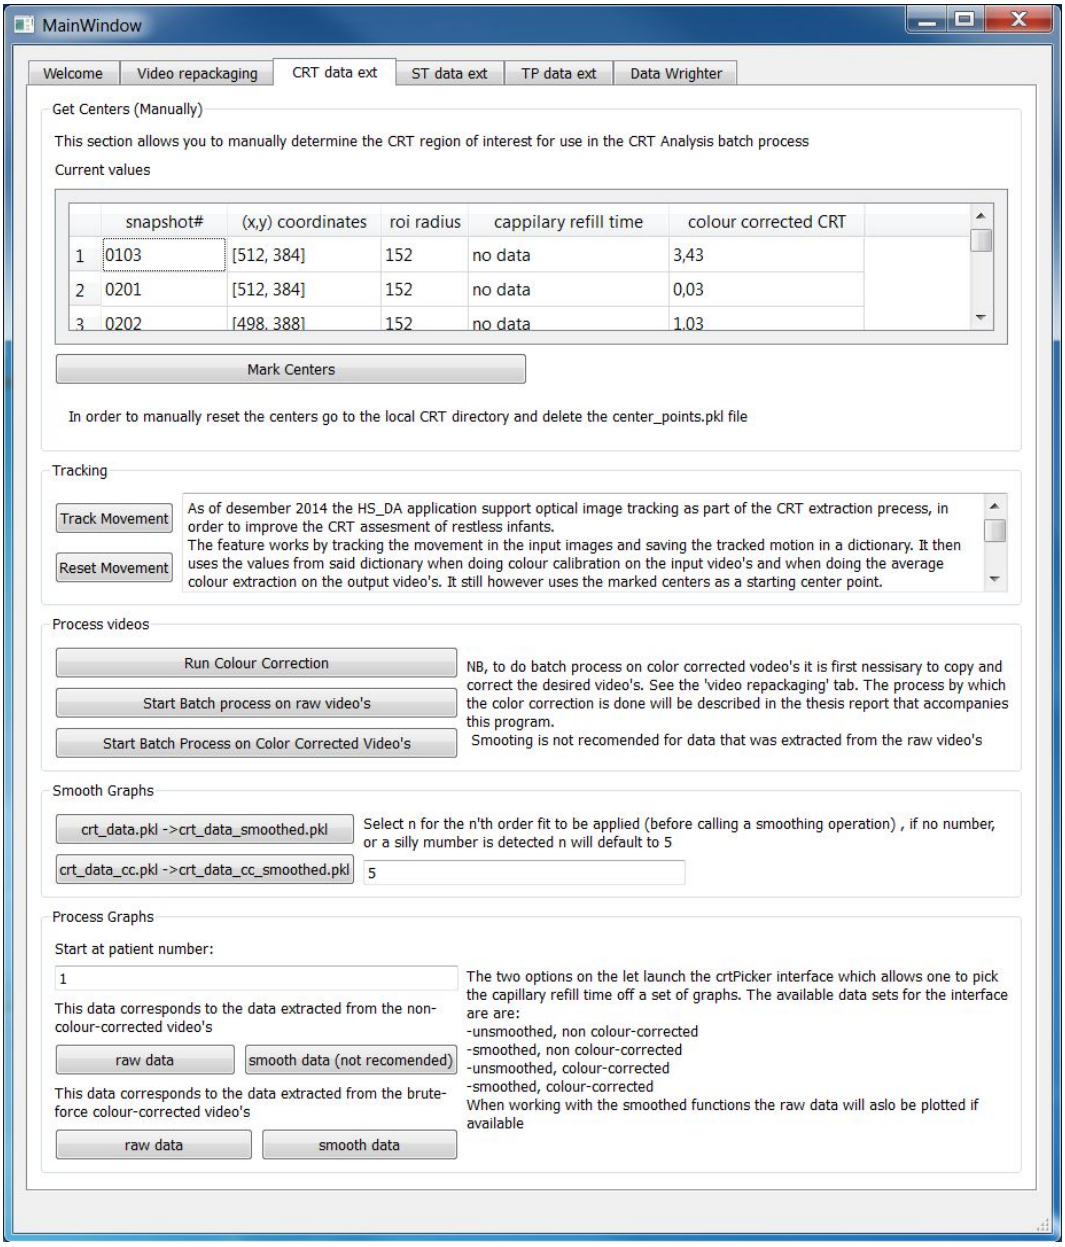

Figure C.2: Data extraction application, CRT data extraction interface

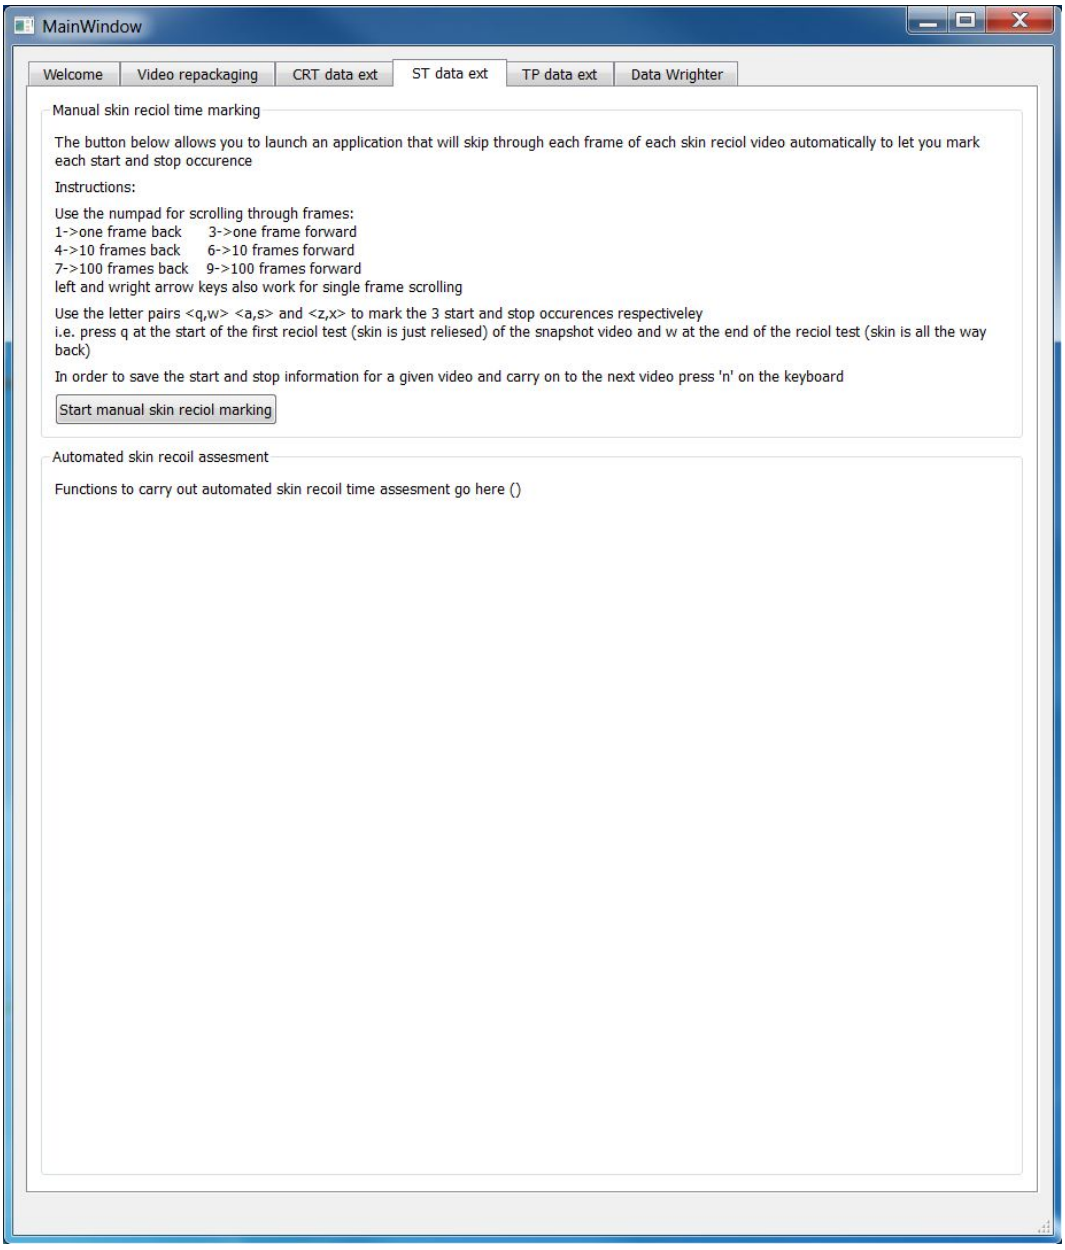

Figure C.3: Data extraction application, SRT data extraction interface

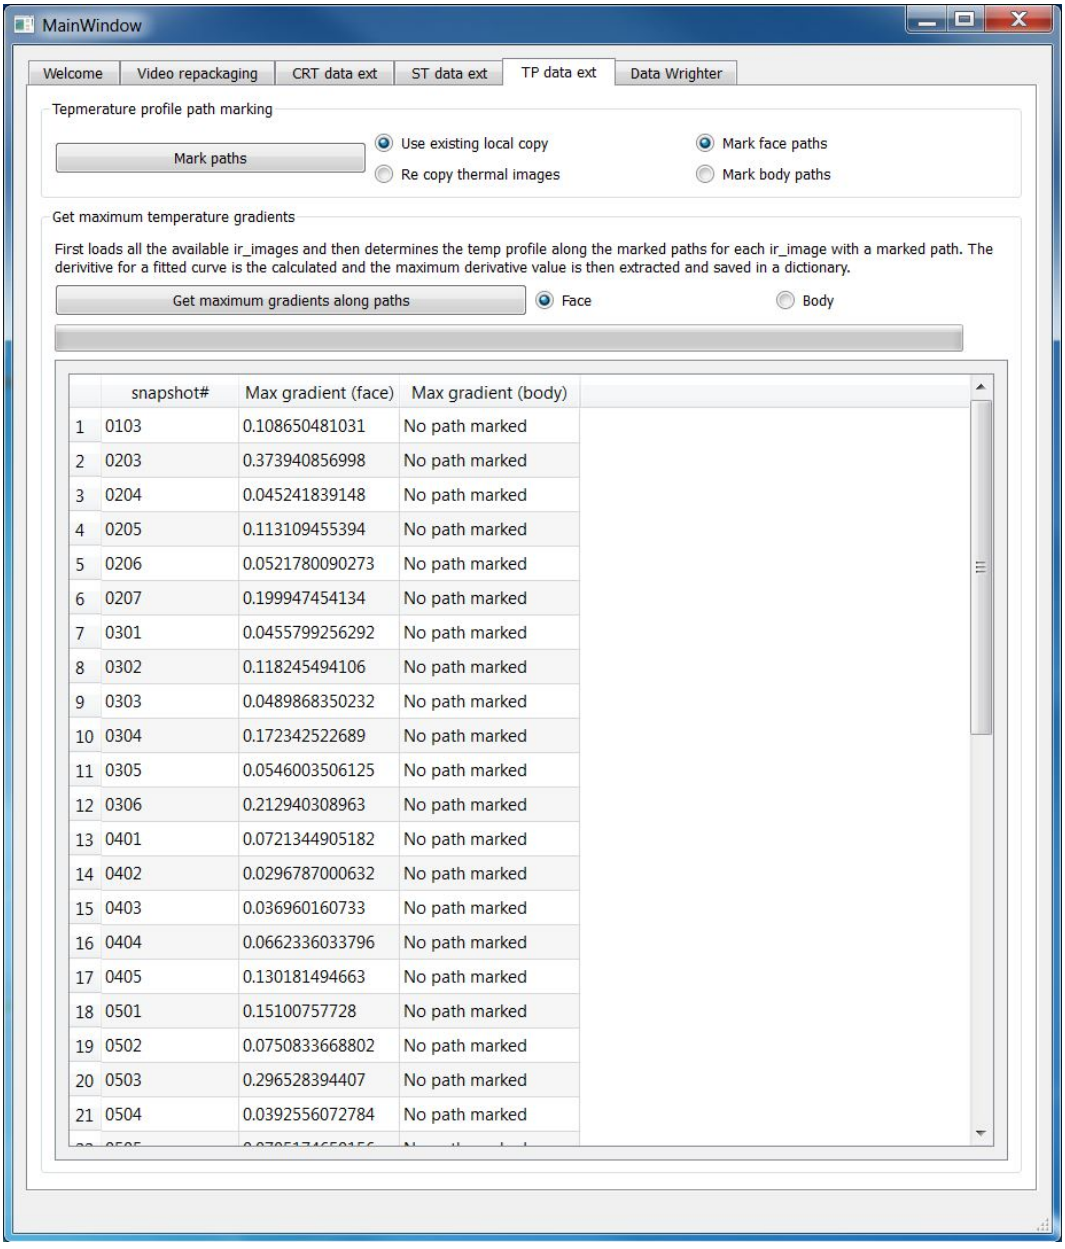

	snapshot#	Max gradient (face)	Max gradient (body)
1	0103	0.108650481031	No path marked
2	0203	0.373940856998	No path marked
3	0204	0.045241839148	No path marked
4	0205	0.113109455394	No path marked
5	0206	0.0521780090273	No path marked
6	0207	0.199947454134	No path marked
7	0301	0.0455799256292	No path marked
8	0302	0.118245494106	No path marked
9	0303	0.0489868350232	No path marked
10	0304	0.172342522689	No path marked
11	0305	0.0546003506125	No path marked
12	0306	0.212940308963	No path marked
13	0401	0.0721344905182	No path marked
14	0402	0.0296787000632	No path marked
15	0403	0.036960160733	No path marked
16	0404	0.0662336033796	No path marked
17	0405	0.130181494663	No path marked
18	0501	0.15100757728	No path marked
19	0502	0.0750833668802	No path marked
20	0503	0.296528394407	No path marked
21	0504	0.0392556072784	No path marked

Figure C.4: Data extraction application, thermal profile data extraction interface

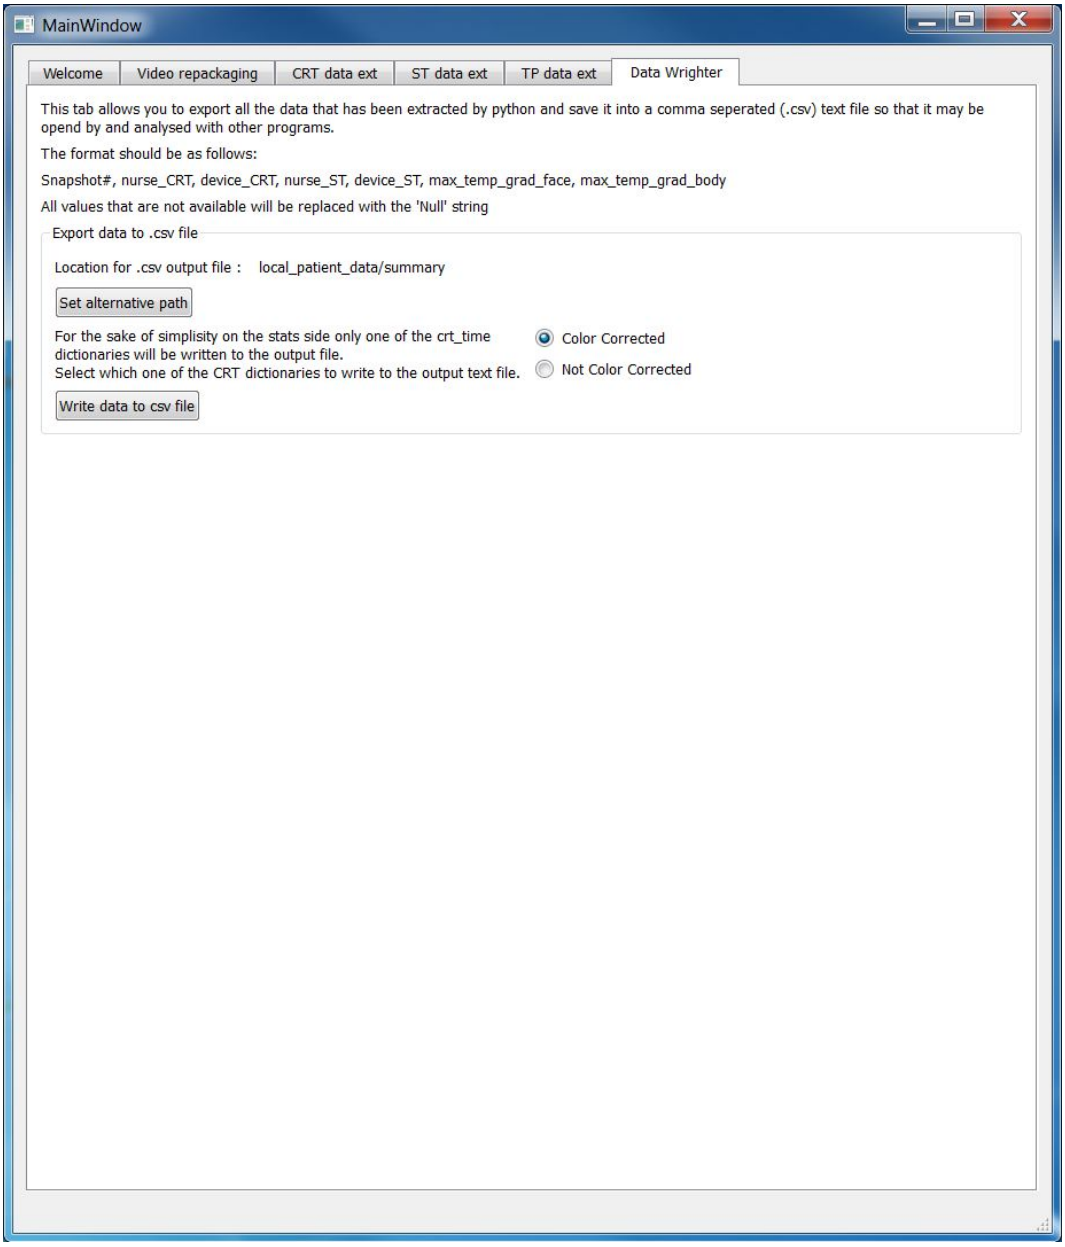

Figure C.5: Data extraction application, data export interface

D. Data sheets

FLIR E-Series

The New E40, E50 and E60, Now with MSX® Advanced Thermal Imaging Camera Performance

Troubleshoot more efficiently, create detailed reports easier, and share images and findings faster with FLIR's latest E-Series thermal imagers. Featuring a fresh array of imaging, communication, and productivity tools to help you get more done in a day.

FLIR Tools Mobile Communication – *Connect an E-Series camera to smartphones and tablets with our Wi-Fi app. Stream live video so others can watch along. Import radiometric JPEGs, adjust levels, add measurement tools, then send images off in simple reports via email to expedite repair decisions.*

Superior Thermal Imaging – *Up to 76,800 pixels (320 x 240) for better long-range accuracy and the highest level of point & shoot camera IR resolution.*

MSX Thermal Image Enhancement – *See numbers, labels and other key visual details not normally apparent in a regular thermal image with an all-in-one thermal picture that instantly orients you to right where heat issues are.*

Onboard Digital Camera – *3.1MP resolution provides clear visible light reference pictures; includes built-in LED lamp that doubles as a flashlight plus a laser pointer to mark locations.*

Large Touchscreen with Auto-Orientation – *The E-Series touchscreen provides an intuitive interface that takes full advantage of the entire 3.5" display for easy access to diagnostics tools, and measurement and other data overlays stay upright no matter how you hold the camera.*

Picture-in-Picture – *Overlay thermal images onto digital scenes for extra location documentation.*

Multiple Measurements – *Add up to 3 box areas and 3 moveable spots using the touchscreen to gather more detailed temperature information.*

MeterLink® – *Measure more than temperature. Send diagnostics data by Bluetooth from compatible clamp and moisture meters directly to E-Series cameras to embed extra information into thermal images as you capture them that further supports findings.*

Interchangeable Lenses and Manual Focus– *Optional wide-angle and telephoto optics let you capture more of a scene in one shot or get in closer to measure smaller, distant targets. Full focus control help you dial in the sharpest clarity and accuracy.*

Higher Temperature Range & Sensitivity – *Measure hotter-running equipment with higher 650°C (1,202°F) detection, and see subtler temperature patterns and detail with E-Series' greater sensitivity.*

FLIR Tools Mobile Wi-Fi Connectivity

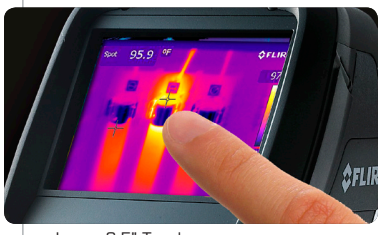

Large 3.5" Touchscreen

Auto-Orientation Keeps Diagnostics Overlays Upright

Visible Light Pictures Align with Thermal Images: 3.1MP Digital Camera, LED Lamp, and Laser Pointer

Large Buttons for Easy Glove-Handed Operation

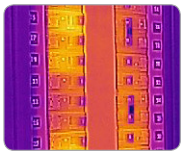

MSX Image Enhancement

Bluetooth® METERLiNK® Communication

Applications

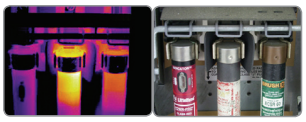
Electrical: Hot Fuses

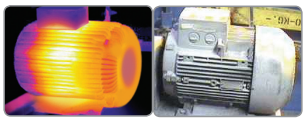
Motor: Internal Winding Problem

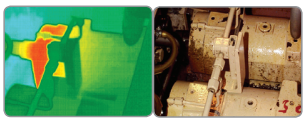
Motor: Bearing Problem

Imaging Specifications

FEATURES	FLIR E40	FLIR E50	FLIR E60
Temperature range	-4 to 1202°F (-20 to 650°C)	-4 to 1202°F (-20 to 650°C)	-4 to 1202°F (-20 to 650°C)
Thermal sensitivity (N.E.T.D)	<0.07°C at 30°C	<0.05°C at 30°C	<0.05°C at 30°C
Detector Type - Focal plane array; (FPA) uncooled microbolometer	160 x 120 pixels	240 x 180 pixels	320 x 240 pixels
MSX® Thermal Image Enhancement	Yes	Yes	Yes
Picture-in-Picture (P-i-P)	Fixed P-i-P	Scalable P-i-P	Scalable P-i-P
MPEG 4 Video Recording	Yes	Yes	Yes
Video Camera w/Lamp & Laser	3.1MP/LED Lamp/Laser pointer	3.1MP/LED Lamp/Laser pointer	3.1MP/LED Lamp/Laser pointer
Digital Zoom	2X Continuous	4X Continuous	4X Continuous
Image annotation	Voice (60s)/Text Comments	Voice (60s)/Text Comments	Voice (60s)/Text Comments
Moveable Spot	3 Spotmeters	3 Spotmeters	3 Spotmeters
Area Box	3 Area Boxes (full image with min/max/avg)	3 Area Boxes (full image with min/max/avg)	3 Area Boxes (full image with min/max/avg)
Delta T	Yes	Yes	Yes
Data Communication Interface	USB-mini, USB-A, Composite Video, Bluetooth, Wi-Fi	USB-mini, USB-A, Composite Video, Bluetooth, Wi-Fi	USB-mini, USB-A, Composite Video, Bluetooth, Wi-Fi
COMMON FEATURES			
Frame Rate	60Hz		
Field of view / Focus	25° x 19° / Manual (Minimum focus distance 1.3ft/0.4m)		
Spectral range	7.5 to 13µm		
Display	Built-in 3.5" color LCD		
Image modes	Thermal, visual, MSX, and image gallery		
Image Storage	>1000 radiometric JPEG images (SD card memory)		
Laser Classification/Type	Class 2/Semiconductor AlGaInP Diode Laser: 1mW/635nm (red)		
Set-up controls	Mode selector, color palettes, configure image info, units, language, date and time formats, and image gallery		
Measurement modes	Auto hot/cold spot, Isotherm (above/below/interval)		
Measurement Correction	Reflected ambient temperature & emissivity correction		
Battery Type/Operating Time	Li-Ion/ >4 hours, Display shows battery status		
Charging system	In camera AC adapter/2 bay charging system		
Shock/Vibration/Drop / Encapsulation; Safety	25G, IEC 60068-2-29/ 2G, IEC 60068-2-6/ Drop-proof 2m (6.6ft) IP54; EN/UL/CSA/PSE 60950-1		
Dimensions/Weight	9.7x3.8x7.2" (246x97x184mm)/<1.82lbs (825g), including battery		

Ordering Information

64501-0101 ..... FLIR E40 IR Camera w/MSX 160 x 120 Resolution/60Hz
64501-0101-KIT ..... FLIR E40 w/MSX 160 x120 Resolution/60Hz with FLIR Tools +
64501-0201 ..... FLIR E50 IR Camera w/MSX 240 x 180 Resolution/60Hz
64501-0201-KIT ..... FLIR E50 w/MSX 240 x 180 Resolution/60Hz with FLIR Tools +
64501-0302 ..... FLIR E60 IR Camera w/MSX 320 x 240 Resolution/60Hz
64501-0302-KIT ..... FLIR E60 w/MSX 320 x 240 Resolution/60Hz with FLIR Tools +

ACCESSORIES
T197752 ..... Li-Ion Rechargeable Battery
T910814 ..... AC Adapter Charger (110-240V, U.S. Plug)
1910490 ..... Cigarette Lighter Adapter Kit, 12VDC (1.2m cable)
T198125 ..... 2-Bay Battery Charger including Power Supply (multi plugs)
T197771 ..... Bluetooth® Headset
T127648 ..... FLIR Tools+ (Scratchcard) 2 Users
1196961 ..... 15° Telephoto Lens
1196960 ..... 45° Wide Angle Lens
T197926 ..... Tripod adapter
T127100 ..... Sunshield
T911087 ..... Pouch

5-10-2 Protection

10-Year Detector Protection
5-Year Battery
2-Year Parts & Labor

FLIR®

BOSTON
FLIR Systems, Inc.
PH: +1 866.477.3687
PH: +1 978.901.8000

PORTLAND
Corporate Headquarters
FLIR Systems, Inc.
PH: +1 866.477.3687

CANADA
FLIR Systems, Ltd.
PH: +1 800.613.0507

MEXICO/LATIN AMERICA
FLIR Systems Brasil
Av. Antonio Bardella
PH: +55 15 3238 8070

www.flir.com
NASDAQ: FLIR

E. Tygerberg forms

Clinical Assesments

Name & Surname			
DOB	/ /		**Folder Number:**
Age		months	
Gender	Male	Female	**Patient Number:**

Date of Admission / /
Time of Admission _____
Weight on Admission _____ kg

History of presenting compaint
Past Medical History

RVD status	Unknown	Unexposed	Exposed	Negative	Positive	
Diet	Exclusive Breast		Exclusive Formula		Mixed/complementary feeds	

Most recent healthy weight plotted on RTHC _____ kg Date of measurement _____

Interpretation of nutritional status based on weight-for-age growth chart in RTHB				
Obese	Moderately Overweight	Normal	Moderately Underweight	Severe Underweight
>+3	>+2	between +2 and -2	<-2	<-3

Consent taken	YES	NO

Clinical assesment

Time: _____ date: _____

For Physicians to complete: *(please tick)*

Sunken Fontanelle	Very	Sunken	None
Sunken eyes	Very	Deepset	None
Mucous membranes	Cracked	Dry	Moist
Neurological status	Depressed / Lethargic	Restless / Irritable	Consolable & Thirsty
CRT	> 3 sec	<3 sec	<2 sec
Skin turgor	> 2 sec	slow but < 2 sec	<2 sec
Tears	Absent	Decreased tears	Normal
Quality of Pulse	Normal	Weak	

Degree of dehydration	SHOCK	Severe	Moderate/Some	None

Blood drawn:	YES	NO

Results:

Clinical assesment

Time: _____ date: _____

For Physicians to complete: *(please tick)*

Sunken Fontanelle	Very	Sunken	None
Sunken eyes	Very	Deepset	None
Mucous membranes	Cracked	Dry	Moist
Neurological status	Depressed / Lethargic	Restless / Irritable	Consolable & Thirsty
CRT	> 3 sec	<3 sec	<2 sec
Skin turgor	> 2 sec	slow but < 2 sec	<2 sec
Tears	Absent	Decreased tears	Normal
Quality of Pulse	Normal	Weak	

Degree of dehydration	SHOCK	Severe	Moderate/Some	None

Blood drawn:	YES	NO

Results:

4 hourly observations

Name & Surname			
Age		months	**Patient Number: 01**
Gender	Male	Female	

Most recent healthy weight plotted on RTHC_____ kg Date of measurement _____

Interpretation of nutritional status based on weight-for-age growth chart in RTHB				
Obese	Moderately Overweight	Normal	Moderately Underweight	Severe Underweight
>+3	>+2	between +2 and -2	<-2	<-3

Consent taken	YES	NO

4 hourly observations

QR code 01-01 

TIME _____ DATE _____

A. For Nursing staff to complete:

Weight _____ kg Tempr _____ C

BP _____ mmHg Pulse rate _____ /min

Respiratory rate _____ /min

INTAKE | IV | oral | both | OUTPUT in last 6 hrs PU x _____

Total intake in las 6 hrs _____ ml BO x _____

B. For own use:

CRT

ST

Termal Profile File Name _____

SC and IR

4 hourly observations

QR code 01-02 

TIME _____ DATE _____

A. For Nursing staff to complete:

Weight _____ kg Tempr _____ C

BP _____ mmHg Pulse rate _____ /min

Respiratory rate _____ /min

INTAKE | IV | oral | both | OUTPUT in last 6 hrs PU x _____

Total intake in las 6 hrs _____ ml BO x _____

B. For own use:

CRT

ST

Termal Profile File Name _____

SC and IR

Parental information leaflet and consent form

Title of the research project: Quantitative hydration sensor development: Infant testing

Investigators:
Mr. Eduard Kieser, Mr. Cobus Visser, Prof. Johan Smith, Prof. C Scheffer
Address:
Mr. E Kieser
Department of Mechanical and Mechatronic Engineering
Room 616
C/O Banhoek & Joubert Streets
Stellenbosch
7600
Contact Details:
Office: (021) 808 3613
Cell: 076 459 6484 (Eduard), 076 907 9121 (Cobus)
E-mail:
cscheffer@sun.ac.za (Scheffer)
js7@sun.ac.za (Smith)
ekieser@sun.ac.za (Kieser)
cobusmetnc@gmail.com (Visser)

Dear Patent/Guardian.

My name is Eduard Kieser and I am a researcher at the Biomedical Engineering Research Group (BERG) at Stellenbosch University, I would like to invite you to participate in a research project that aims to investigate how dehydration works in babies.

Please take some time to read the information presented here, which will explain the details of this project. Please ask the study staff or doctor any questions about any part of this project that you do not fully understand. It is very important that you are fully satisfied that you clearly understand how this research affects you. Also, your participation is entirely voluntary and you are free to decline to participate. If you say no, this will not affect you negatively in any way whatsoever. You are also free to withdraw from the study at any point, even if you do agree to take part.

This study has been approved by the Health Research Ethics Committee (HREC) at Stellenbosch University and will be conducted according to accepted and applicable National and International ethical guidelines and principles, including those of the international Declaration of Helsinki October 2008.

What is the study all about?

Dehydration in babies is difficult to detect. Healthcare workers in poor areas often lack the experience required to make an accurate hydration assessment. The purpose of this study is to test how well a set of sensors that are being developed work. The study will be done at the Tygerberg Hospital, where 40 patients will be observed during their treatment.

Blood flow test

The first device will test the baby's blood flow to the skin. The device is made up of a camera and a pressure pad that is used to apply a small pressure on the baby's chest. The device will be placed on the chest, the pressure pad shown below will be used to apply a small pressure for between 3 and 7 seconds. The pressure pad will then flip out of the way to allow the camera to start recording to see how long it takes for the colour to return to the part where the pressure was applied.

Skin recoil test
Skin recoil time is also commonly used to check for dehydration. To manually measure the skin recoil of a baby, the doctor would normally pinch and pull on the skin of the stomach, let it go and measure how long it takes for the skin to snap back. To make the test more accurate a device was developed to digitally measure the recoil time. The device consists of a camera to record the skin recoil frame by frame and store it for later analysis. For the test a line will be drawn, with a medical marker, right above the belly-button from side to side. The skin recoil camera will then be placed on the chest looking at the stomach and the line that was drawn on it. The doctor will then test the skin recoil at the stomach the way he/she normally would while the camera records the test. The video will then be slowed down to measure the recoil time more accurately

Skin temperature profile
When patients are dehydrated, less blood flows to the skin which means that the patient's surface temperature is less even. During this study a thermal (heat sensing) camera will be used to capture the patient's temperature at various times during his/her recovery to see if there is a link between the temperature profile and the patient's hydration state.

Skin hydration probe
This test will involve placing a dry electrode on the arm, stomach and legs of the patient. The electrode will be kept against the skin for less than a minute.

Light dehydration sensor
For this test a small light and sensor will be gently pushed against the skin. The light will be very weak and similar to the light that comes from a television remote (almost invisible).

Why have you been invited to participate?
For the study we need to measure the symptoms of dehydrated patients. Diarrhoea is one of the most common causes of dehydration in infants. Your son/daughter was admitted to Tygerberg with diarrhoea and expected dehydration. During the course of the study we will aim to test 40 patients.

What are my responsibilities?
You do not have any responsibilities. If you agree to participate in the study your child will receive the normal treatment with the exception of the procedures explained in this document. Your child will also not be kept at the hospital for longer than is required for his/her treatment.

Will you benefit from taking part in this research?
Your child will not directly benefit from this project, since the results of the study tests will not be used in the treatment of your child.

Are there any risks involved in taking part in this research?
None of the devices that will be used for the study pose any reasonable threat. The biggest expected risks are associated with handling and moving of the patient if and when necessary.

If you do not agree to take part, what other alternatives do you have?
As mentioned before all participation in the study is entirely voluntary. Any participant can at any stage withdraw from the study, without affecting the health care they will receive in any way.

Who will have access to your medical records?
All information collected will be treated as confidential and protected. If it is used in a publication or thesis, the identity of the participant will remain anonymous. Only HREC members, the study investigators and sponsors will have direct access to the raw information collected.

Will you be paid to take part in this study and are there any costs involved?
No, you will not be paid to take part in the study. Participation in the study will also not cost you anything.

What if something goes wrong?
In the unlikely event that your child is hurt as a result of his/her participation in the study, the supervising physician Prof. Johan Smith will be informed immediately and he will take the appropriate steps.

Is there anything else that I should know or do?
You can contact Prof. Smith at (021) 938 4835 or myself at 076 459 6484 if you have any further questions or encounter any problems. You can contact the Committee for Human Research at (021) 938 9207 if you have any concerns or complaints that have not been adequately addressed by your study doctor.

You will receive a copy of this information and consent form for your own records.

Declaration by Participant

By signing below, I agree to take part in a research study entitled "Quantitative hydration sensor development: Infant testing".

I declare that:

1. I have read or have been read to the contents of this document and agree to all information and/or conditions given therein.
2. Ample time has been given to me to ask and/or question the information and/or conditions in this document.
3. All questions that were posed by myself have been answered and I fully understand everything in this study
4. I understand the risk involved in this study.
5. I understand that taking part in this study is voluntary and I have not been pressurized to take part
6. I may choose to leave the study at any time and will not be penalized or prejudiced in any way.
7. I acknowledged that I am a volunteer participating in this study out of free will and have not been forced to do so.

Signed at ____:____ on ____/____/2014

Signature of Participant_____

Signature of Witness_____

Declaration by Investigator

I, (the investigator) , hereby declare that:

1. I have fully explained the information in this document to the guardian of the participant.
2. I have informed the applicant to the best of my abilities.
3. I have encouraged the applicant to ask about and/or question the risks involved in this study.
4. I have encouraged the applicant to ask about and/or question the information and conditions set out in this document.
5. That I have not forced the applicant to volunteer for the study.

Signed at ____:____ on ____/____/2014

Signature of Investigator _____

Signature of Witness _____