Motion Artefact Reduction for Reflection-Mode Photoplethysmography

Matthew James Butler, MEng

Optics and Photonics Group,

Department of Electrical and Electronic Engineering,

University of Nottingham

Thesis submitted to the University of Nottingham

for the degree of Doctor of Philosophy

February 2018

Abstract

Photoplethysmography (PPG) is a technique that uses light to measure the local changes in blood-volume in subjects (predominantly humans). Multiple useful statistics can be gained from such a measurement; heart-rate and it's variability, blood-oxygen saturation and even an estimation of blood pressure, to name but a few. Compared to other measurement techniques, photoplethysmography is favourable as it is both non-invasive, since nothing physical penetrates the subject's skin, and safe, as the subject is galvanically isolated from the test equipment (additional benefits also exist).

Motion artefacts (errors in the measured signal caused by physical movement) are the largest source of error when photoplethysmographic measurements are made, and with the majority of applications involving some form of movement, a motion-tolerant PPG extraction technique would allow for more precise recordings/research/diagnosis etc. This thesis presents the development of an improved photoplethysmography technique that has increased resilience to motion.

The developed technique uses multiple PPG measurements at different locations to reconstruct a single PPG signal. It is shown that despite the signals being taken in close proximity to each other (less than 3 cm separation between the farthest elements), the variation in the signals gives sufficient redundancy to extract the uncorrupted PPG to a much higher accuracy using Independent Component Analysis, achieving in the worst case, a 78% reduction in the calculated artefact presence (using quality calculating functions, also presented).

As the vast majority of existing PPG systems use a single sensing element, it is hypothesised that such systems cannot be used to accurately

and continuously detect the PPG for most motion types and severities. A working prototype of the developed system is demonstrated and directly compared to a single-channel system, showing its effectiveness.

Acknowledgements

I would like to express my sincere gratitude to my supervisors, Prof. Barrie Hayes-Gill, Prof. John Crowe, and Dr. Paul Rodmell, for their unfaltering support, technical input, and guidance, throughout this research project.

Additionally, I would like to thank my examiners, Dr. John Walker (internal) and Dr. Justin Phillips (City, University of London) for their comments and critiques that enabled me to improve and finalise this thesis.

This research has been funded by the Engineering and Physical Sciences Research Council (EPSRC, grant number EP/K503101/1), and supported via a Cooperative Awards in Science and Engineering (CASE) scheme with Tioga Ltd (Derby, UK), both to which I extend my thanks.

Finally, a huge thanks to my family and friends, for their continued support and encouragement.

Statement of Originality

To the best of my knowledge, the contents of this thesis, as presented, represent my own work.

The two pieces of hardware and associated firmware/software presented in this thesis (the PPG/ECG Recorder and Optical Matrix as shown in Chapters 4 and 5, respectively) are of my own design and construction. The top-level design of the algorithms presented in Chapter 4 (Heart-Rate and Quality Level estimation) are of my own design, using (where stated/cited) already existing and third-party, lower-level functions (FFT, etc).

This thesis has not been submitted for any other degree or purpose. Some of the results and outcomes acquired in Chapter 3 have been published in a journal paper (see details within the chapter).

Matthew James Butler

Contents

Chapter 1

Introduction

Every physical quantity will have its own variety of associated measurement techniques. These techniques vary in terms of the required hardware, processing methods and perhaps a preferred mode of displaying the results. All will be prone to different errors; both due to inaccuracies in the technique itself, and due to external influences when the measurement takes place.

Some of these external influences are easy to remove and can be attenuated by simply moving further from them, or via filtering. Others however, are not so easy to remove, and can corrupt the desired signal entirely, or at the very least, introduce artefacts.

For certain applications, some artefacts may be tolerable, and obtaining a 'correct' result even just for a small fraction of the time may be acceptable. However, for applications where consistent and reliable measurements are required - some medical analyses, for example - finding ways to suppress the errors is not only desirable but a necessity.

1.1 Photoplethysmography

The terms photoplethysmography and photoplethysmogram (both acronymised as PPG) derive from the Greek 'photos' *light*, 'plethys' *fullness* and 'graphos' *to write*: or in modern vernacular, "measuring volume using light". More specifically, in a medical context, the volume most often refers to that of pulsatile blood [1].

1.1.1 Pulsatile Blood-Flow

Blood within veins, arteries and the interconnecting capillaries is pulsatile due to the rhythmic contractions of the heart, the effects of which extend throughout the whole body. These rhythmic contractions are known as the *cardiac cycle* [2].

1.1.1.1 Cardiac Cycle

For a healthy human, the heart will 'beat' periodically, pumping blood around the body. Two main circulatory systems exist (there are more sub-systems), each emanating from and then returning to the heart (see Figure 1.1).

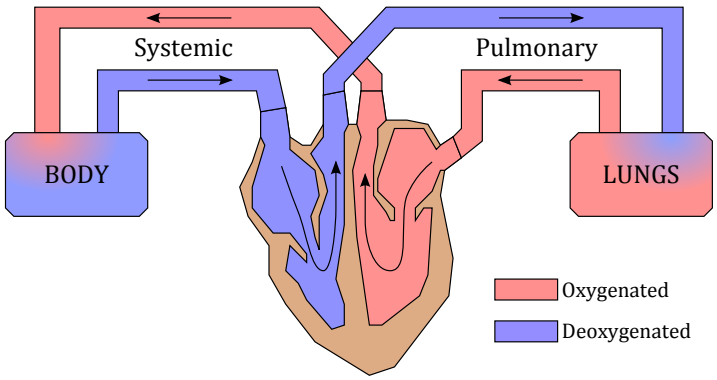

Figure 1.1: *Illustration of the heart, and the two primary circulation systems.*

The first, known as the Pulmonary Circulation, transports blood through the lungs. Here, blood which has been depleted of oxygen and carrying waste carbon dioxide is rejuvenated by trading the gases with inhaled air. Blood which returns to the heart is enriched with oxygen and no longer has the burden of the carbon dioxide. This system is contained entirely within the chest, and is not easily accessible without invasive procedures.

The second system, known as the Systemic Circulation, transports this oxygen-enriched blood around the rest of the body. As parts of this system pass close to the surface of the skin (close enough to be visible without any

artificial aids), the Systemic Circulation is easily accessible for measurement and is the focus of study for this thesis.

Each heart 'beat' can be simplified to include just two periods: *diastole* and *systole*. During the diastolic phase, the muscles in the heart relax allowing blood to flow into it. Pressure in the Systemic Circulation *generally* decreases down to its minimum as the blood is not under compression from the heart (a small caveat to this is presented in Section 1.1.1.3).

A contraction of the muscles in the heart, increasing pressure in the Systemic Circulation, happens in the systolic phase. Pressure in the main artery leaving the heart (towards the Systemic Circulation: Aortic Pressure) is plotted in the "Wiggers Diagram" [3] in Figure 1.2 as line 'AP', the shape of which is important (discussed later).

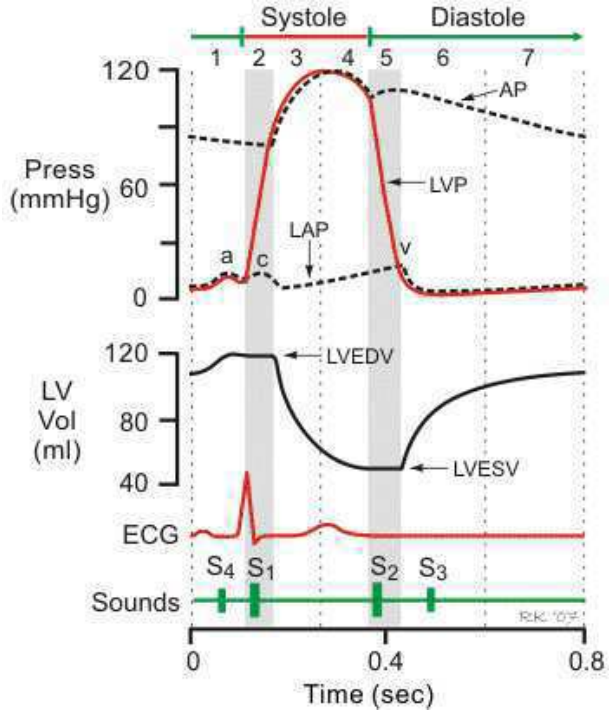

Figure 1.2: *Pressure and volume of regions of the heart during one cycle. The aortic pressure (AP) is the most relevant when discussing photoplethysmography. (Image unedited from [3]).*

1.1.1.2 The Heart's Effect on the Body

As the arteries from the heart carry higher pressure blood, and as all blood vessels have some elasticity, their mechanical expansion can be easily felt with just the fingers (this is the palpation of the well-known 'pulse') [4]. However, as arteries lie deep within the tissue, they are not readily accessible and their mechanical effects can only be felt at certain locations of the body where they run closer to the surface; most notably at the wrists and neck [5–7].

Contrarily, veins carry lower pressure blood and appear closer to the surface of the skin. By the time the blood has travelled through the arterial and capillary networks to the veins, the pressure has dropped and there is seldom any mechanical effect. Placing a finger over a vein, for example, will simply restrict (and probably stop) the blood flow. By blocking the flow at one point and 'scraping' another finger along a vein in the direction of the blood-flow, effectively blanching it, it is possible to see the speed of the blood at that location as it re-fills the vessel (see Figure 1.3).

As arteries branch into smaller vessels (arterioles), both the blood pressure and speed decreases. The finest branches - capillaries - despite their small size compared to other vessels, are more numerous and occupy a significantly larger volume of the skin and tissue. Capillaries are present in much of the tissue around the body and have the role of delivering oxygen enriched blood, and exchanging nutrients for waste products in their surroundings. Much in the same way as the veins, if blanched through compression then released, blood will return slowly; however, whereas with the veins re-filling by 'flowing' from the previously blocked point, the vessels will do so from beneath in a less ordered fashion resulting in a uniform colour change (see Figure 1.4).

Just as with the veins, capillaries exert little mechanical effect on the surroundings. Only when a region with high capillary and nerve ending

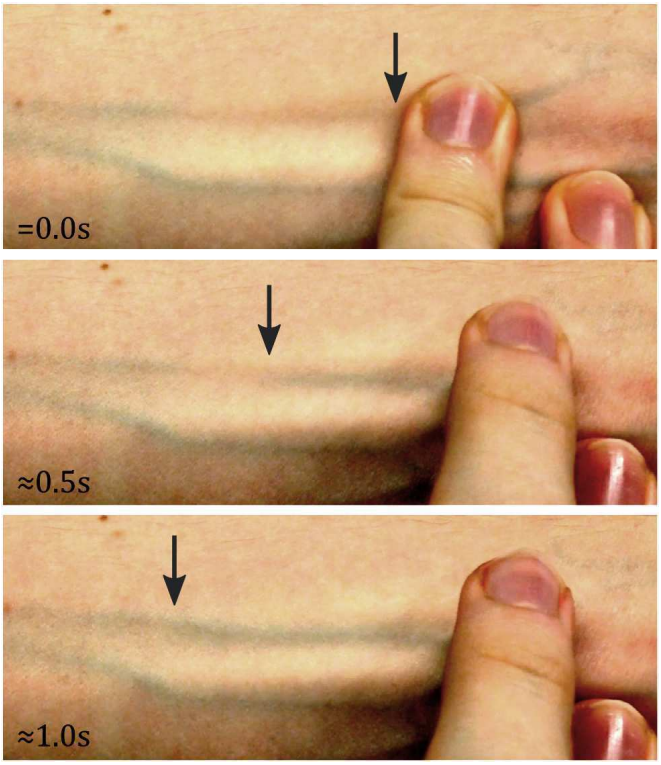

Figure 1.3: *Demonstration of the low pressure and velocity of blood flow in surface veins. The arrows show the position of the venous blood's 'wave-front' as it re-fills the vein.*

densities (making the area more sensitive to touch) is compressed is it possible to 'feel' the volume change caused the the accumulated blood pressure (due to the heart-beat) in the capillary bed. One such area is the tip of the thumb, which is why it is often said that the thumb should not be used for taking someone else's pulse [8].

An increase in the internal (blood) pressure of the source arteries has been reported to cause local capillary density to similarly increase, this is due to the compression of the surrounding tissue [9]. When blood is forced through the capillaries, whose internal size is not that different from that of a red blood cell, and hence having a higher resistance, the cells are also squashed slightly closer together [10]. The increased pressure from a heart-beat will ultimately cause a local increase of red blood cell density in the surrounding tissue.

The local blood pressure (which is modulated as a result of the heart's

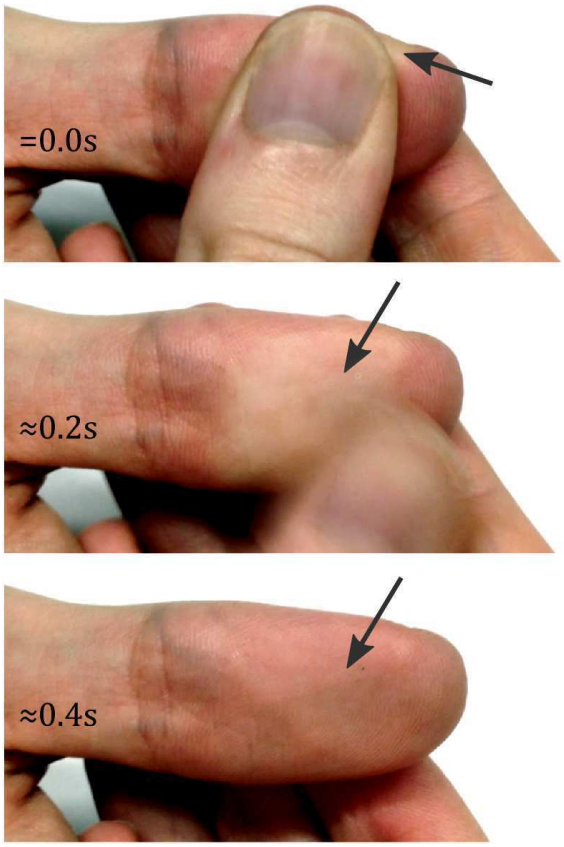

Figure 1.4: *Capillary refill showing a (generally) uniform colour change over the surface of the skin.*

cycle) and the volume of blood in a particular region are related, measuring the volume ultimately allows for information to be obtained about the heart. The modulated blood pressure will be visible as deviations in the detected blood volume; from this, for example, the heart-rate can be extracted.

1.1.1.3 Dicrotic Notch

Although the aortic pressure plotted in Figure 1.2 ('AP') shows an increase during the systole phase and a decrease in the diastole, there is a small irregularity at the start of the latter (highlighted in region '5' in the figure) whereby the pressure momentarily increases. This increase is due to the elasticity of the aorta after the aortic valve (at the junction between the heart and the aorta) closes [11]; i.e. the closure of the valve causes a

temporarily hastened reduction in pressure due to the reversal of blood flow, then after the valve has closed completely, the pressure increases as the artery wall's elasticity tries to compress the restricted blood. Despite the relatively small size of the valve and volume of blood in the region, this pressure change can be detected throughout the body.

1.1.2 Blood/Tissue Optical Characteristics

There are competing theories and models for what causes the cardiac-synchronous variations in light that is returned to the skin's surface after illumination [9, 12–14]. Some models look at the absorption of light by varying geometry, some consider variations of the effects of scattering. Either way, the increase in blood pressure caused by each heart-beat has a mechanical effect on the body, which in turn affects how light from the surface progresses.

When light travels through a medium, two main effects are of relevance when discussing PPG: absorption and scattering. Absorption is more intuitive as the effect is easily observed and understood; the energy contained in the electro-magnetic radiation (light) is transformed to heat energy within the absorbing matter. The light effectively ceases to exist and is no longer part of the system. Scattering (of light), on the other hand, comes in many types (Rayleigh, Mie, Raman, etc.) and is less intuitive. Put simply, light scatters when interacting with matter in its path and causes it to, at the very least, change direction. Materials that have low scattering 'coefficients' appear transparent (air, glass) as the light passes through with minimal interaction, whilst materials with higher coefficients of scattering appear more diffuse (paper, skin/tissue, etc).

It is the optical characteristics of blood and the skin which make the non-invasive measurement of pulsatile blood possible.

1.1.2.1 Blood

Blood's absorption of light is of particular interest as using light to measure its volume, as previously mentioned, allows for a convenient and non-invasive method of determining heart-rate, for example. Additionally, although not looked at in depth in this thesis, the oxygenation level of the blood can be determined by using the fact that its absorption spectra changes with oxygenation level (Figure 1.5). A little under a half (by volume) of the blood is composed of red-blood-cells, which contain a protein called haemoglobin [15], responsible for this absorption.

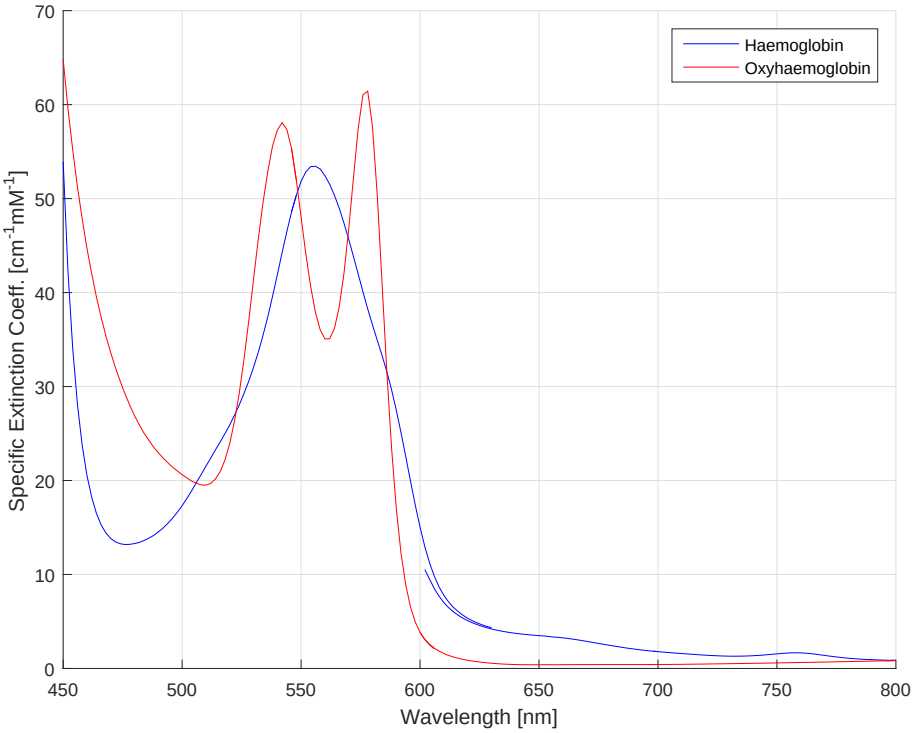

Figure 1.5: *Graph of blood haemoglobin and oxyhaemoglobin absorption at different wavelengths. (Compiled from two data sets from [15]).*

In Haemoglobin's 'de-oxygenated' state, as it would be when returning to the heart through the veins, it is known as *de-oxygenated haemoglobin*. When it has oxygen bound to it, as it would be after passing through the lungs, it is known as *oxygenated haemoglobin*, or more simply: oxyhaemoglobin. The difference in absorption spectra between these two states is caused by a change in structure due to the bound oxygen.

1.1.2.2 Skin/Tissue

Optical effects of the matter surrounding the blood vessels must be taken into account when considering optical measurement of the blood. Photons that penetrate the skin's surface will either be absorbed or scattered by this tissue. A large optically-absorbent constituent of the skin is Melanin [16], the absorption spectrum of which is presented in Figure 1.6. This defines the skin colour.

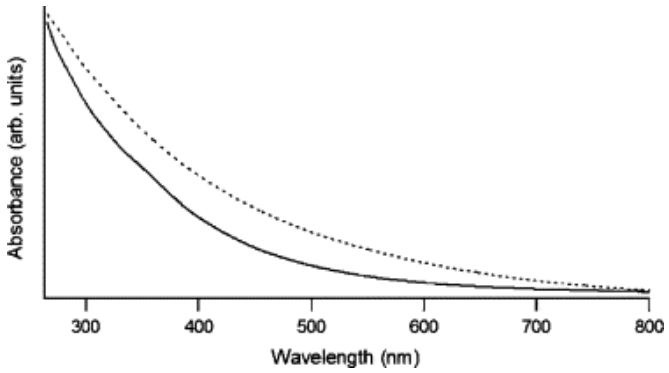

Figure 1.6: *Graph of melanin optical absorption vs wavelength (two subtypes: brown-black eumelanin (dashed line) and brown-red pheomelanin (solid line)). (Image from [17]).*

Light that enters the skin at a single point will appear to leave at different locations with decreasing intensity as the distance from the entry point increases (Figure 1.7). This is due to two reasons, firstly, the longer the optical path is, the more likely the photons will be absorbed. And secondly, as the light travels further from the source, the area/volume of material the light is traversing is getting larger, hence higher dispersal (or in other words, the same amount of light in a larger volume results in a lower photon 'density').

As light travels through the skin in this manner, any variation in optical path length between a light source and sensor will likewise cause a variation in the received signal. This is especially problematic when the tissue being measured is deformable, and the desired signal is the much smaller variation in light intensity (via absorption) caused by the pulsatile blood flow.

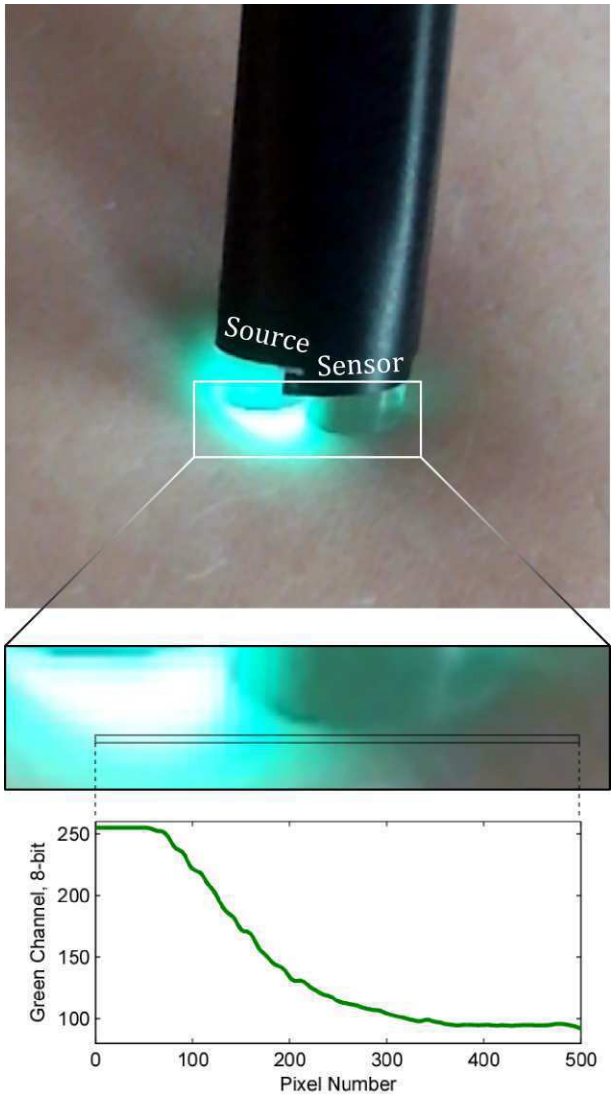

Figure 1.7: *With highly scattering materials such as skin, the photons are more likely to leave closer to the entry point.*

1.2 Methods of Extraction

Measurement of light that has interacted with pulsatile blood (light that has been scattered beneath, and then returned to, the surface) is the existing technique known as Photoplethysmography and has four main modes of operation: in-contact transmission; in-contact reflection; non-contact transmission; and non-contact reflection. Each mode 'quadrant' fundamentally relies on the same physical interaction between light and the blood-/tissue, but each has differing methods of mechanical interaction with the subject, and different additional sources of artefact.

1.2.1 In-Contact

As the name suggests, in-contact photoplethysmography has the probe connected to the subject with no mechanical separation (see Figure 1.8 for examples). In terms of engineering complexity, these are the most basic sensor designs as they often have just one sensing element and no appreciable 'optics' (lenses etc.). This makes them cheap, compact, simple to manufacture, and often when used in a clinical environment: disposable [18].

Figure 1.8: *In-contact sensor arrangements. Images (left to right) from [19–21].*

Because of the optical elements' close proximity to the skin, there is the opportunity to design the sensors such that little or no light 'leaks' from the

source to sensor without passing through the subject. This allows for more efficient designs; also contributing to the efficiency is the fact that little light is 'lost' as the majority enters the skin. For non-contact systems, the detection region would most likely be the same location as where the light enters (due to a light source 'flooding' the region with light), this would result in large reflections from the surface of the skin increasing the chance for artefacts to interfere, and reducing the dynamic range of the system.

As photoplethysmography uses light to detect the pulsatile blood, the power consumption is greater than that of an electrocardiography solution (which is simply *detecting* an electrical signal). Although battery powered solutions do exist (smart watches, for example [22]), their operating times are limited. For this reason the vast majority of clinically rated devices use cables to connect the sensor head, a finger clip for example, to a powered base unit (usually mains powered with a large backup battery). The inconvenience of having cables between the sensor (patient) and the base unit makes the units limited to stationary monitoring, this is not necessarily a problem in a clinical setting. What is a problem, however, is the physical connection itself as this allows the transmission of bacteria etc. requiring expensive and time-consuming cleaning; this may have been a motivation to reduce the size (and especially the cost, to increase profits) of the sensor heads in order to make them disposable.

1.2.2 Non-Contact

In order for optical sensing elements to work correctly at long-range, additional elements such as lenses are required to focus the light. An imaging sensor (camera) is a device that performs this task, is readily available, and is relatively inexpensive (although still more expensive than the in-contact probes, and certainly not disposable). Due to a camera's existing use as an imaging device, the sensing component contains not one element, but

a large array, potentially millions of 'pixels' in size (see Figure 1.9). This gives the potential for advantages such as synchronously detecting pulsatile signals at multiple locations (or from multiple people [23]), which would not be feasible for in-contact set-ups. For some applications where the blood flow over an area of skin needs to be monitored, for example, a camera can be used to take the measurements in one operation; whereas a single 'point' sensor would have to be repositioned for each reading - a laborious prospect.

Figure 1.9: *Non-contact sensor arrangements. Cameras are readily available as non-contact optical sensors. Images (left to right) from [24, 25] and [26] (Fig. 4).*

The lack of physical contact between the sensing element(s) and the patient immediately provides an advantage in a clinical setting as, unlike the in-contact methods, no bacteria or infection can be transmitted. Also, with no mechanical linkage, the subject (patient or otherwise) is free to move unimpeded. Several applications exist that benefit from this freedom (see Section 2.2 for examples).

1.2.3 Transmission Mode

Optical transmission is when light travels through an object after penetrating its surface. For transmission mode photoplethysmography, this suggests that the light source and sensing elements be placed on opposing sides of the subject as in Figure 1.10. Although this is true, it is a slight misnomer as the primary optical interaction with the tissue is, like all modes of PPG, scattering.

So as not to require large - potentially dangerous - light intensities, the optical path from source to sensor must be short enough so that enough light can make it through to be measurable. For this reason, the preferred sites for this mode of operation are limited to small extremities such as fingers, toes, ear-lobes, etc. If a source and sensor were to be placed on opposite sides of a larger body part, a human torso (front to back) for example, then an infeasibly large amount of light[1] would be required for the sensor to detect the transmitted (forward-scattered) light.

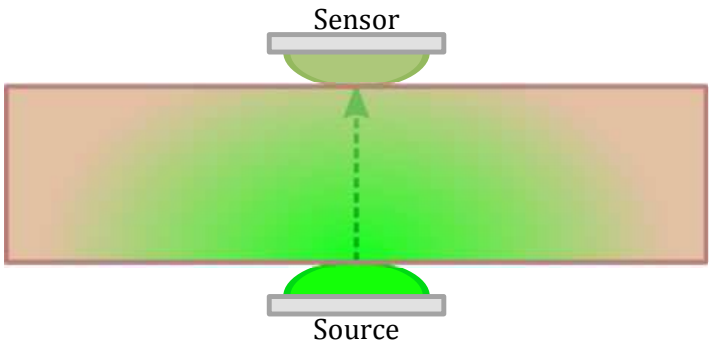

Figure 1.10: *An example in-contact transmission mode (as is used in 'finger clips') showing the light path between the source and sensor.*

One important feature of transmission mode is that whilst the (typical) length of the optical path is dependent on the deformation of the extremity, the tissue that the light is passing through will not change its composition. In other words, regardless of the force applied to the body part, the light will generally always pass through the same components: skin, tissue, blood vessels, bone, more tissue, blood vessels, and through the skin on the other side. Compression of tissue, however, often blanches the local blood vessels, reducing the pulsatile absorbing component [27].

[1]Light in this thesis refers to wavelengths between around 400nm and 1000nm (visible-violet to infra-red). Other wavelengths, such as ≈1nm (X-rays) could pass through tissue of these depths mostly unimpeded.

1.2.4 Reflection Mode

For photoplethysmography measurements, when light entering the surface of the skin scatters back to be detected on the same side, it is known as 'reflection-mode' (see Figure 1.11). This is a slight misnomer as no *reflection* occurs in the traditional sense (primarily); this technique only works because of the scattering properties of tissue.

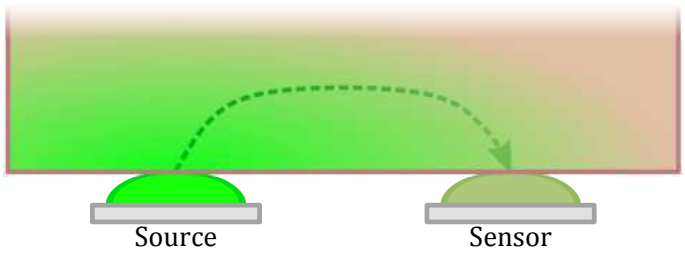

Figure 1.11: *In-contact reflection mode.*

Unlike in transmission mode, the typical optical path length is not generally determined by the topology of the subject (although due to the non-homogeneous structure of the skin, there will be some variation); instead, the mechanical design of the sensor and the wavelength of the light source(s) will influence the optical behaviour. However, counter-intuitively, despite this 'rigid' optical path *length*, the sensitivity of the system to motion is often considerably higher. This is mainly due to the compression of the skin altering the structures that pass through the optical path, including blanching of the vessels (Figure 1.12).

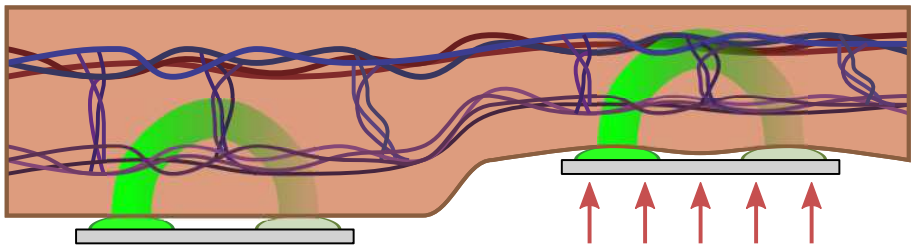

Figure 1.12: *Stylised illustration of light 'paths' with and without compression. When compression (deformation) exists, the light's 'path' will interact with different structures, causing variations in scattering and absorption.*

The major advantage of a reflection mode system is the ability to position the source-sensor probe at any location on the body (assuming pulsatile blood flows beneath the surface).

1.3 Discussion

Due to the absorption profile of the pulsatile blood near the surface of the skin (within capillaries, primarily), wavelengths around 550 nm (green) are particularly suitable for *ex vivo* measurements. This is due to light being absorbed strongly by haemoglobin within the red blood cells, whilst being fairly easily scattered through (and not absorbed by) the surrounding tissue. Other wavelengths, such as those nearer the blue-end of the optical spectrum are far too easily absorbed by the skin that little light would penetrate and scatter back to awaiting sensors. Longer wavelengths (nearer red and infra-red), due to their lower absorption by the skin, have the opposite problem where they would scatter easily through the skin, but be absorbed less by the pulsatile blood[2].

As in-contact transmission-mode PPG acquisition is a very mature field, and is the most widely used in the medical industry (finger/ear clips, for example), and that non-contact (transmission) variants have few benefits, the focus of this thesis will be on reflection-mode photoplethysmography.

[2]Due to the differing absorption characteristics of oxy- and de-oxygenated haemoglobin at these wavelengths, red and infra-red are predominantly used for blood-oxygen-saturation measurements (SPO2).

Chapter 2

Background

Many methods exist that are able to extract the heart-rate from a human. The earliest examples have been to simply count the palpations caused by the heart-beat in a given time - a technique that is still used today [4, 28]. However, as needs for continuous monitoring have increased, automation has been required. From equipment in use in hospitals to consumer level devices that can be incorporated into everyday items, techniques that detect the pulsatile nature of the heart need to become more complex as the activities that the subjects undergo become more energetic (resuscitation in hospitals [29], fitness activities [30, 31], etc). With motion being the primary cause of measurement error, research into methods to reduce its effect are numerous, but none are, as yet, completely effective.

2.1 Brief History

Photoelectric Plethysmography was proposed as a method of measuring pulsatile blood as early as 1937 [32]. Hertzman's hardware set-up used the aforementioned 'transmission-mode' method on far-extremities of the human body (fingers and toes). At this point, no attention was given to optimising the performance of such a system; this was most likely due to the limitations of the technology at the time. A broad-spectrum light source was used with no optical filtering, and the device (combined light source and sensor) was cumbersome and lacked portability. However, it triggered further research in the field to improve the technique.

Even earlier examples of the use of photoelectric plethysmography (later shortened to simply photoplethysmography) exist, however their use was not for measuring the pulsatile blood within subjects as it is primarily used today. Instead, the techniques were perhaps a more literal version of the term 'photoplethysmography' as it was the non-pulsatile blood *volume* that was the goal for measurement. Specifically, the effects that the application of drugs/medication had on the muscle and tissue were the main use, as the volume of blood within the region could be used as an indication of the blood circulation which in-turn was influenced by the injected chemicals [33, 34].

2.1.1 Measurement Artefacts

In Molitor and Kniazuk's (1936) article [34], they mentioned problems that were encountered when performing their experiments regarding errors in their signals: *"...while the method is bloodless in itself, some motor or sensory depressant must be given because even the slightest muscular contractions produce plethysmographic changes..."* [34]. This is the first observed reference to a 'motion artefact', and the extent that such an artefact can affect a *pulsatile* PPG signal can be inferred from their results. Figure 2.1 shows the kymograph of the PPG signal output (along with the temperature) of a rabbit's ear. Accompanying text in the article details the kymograph's specifications, specifically, that the total vertical travel of the needle is 11 cm, with a maximum traversal speed of $8\,\mathrm{cm\,s^{-1}}$. This along with the time scale along the bottom is enough to estimate the effect that motion they are referring to can corrupt the desired signal. Additionally, although not mentioned, it is assumed that the vertical extent of the PPG trace is no less than half the maximum swing.

Given that the detection of a pulsatile PPG is possible with rabbit ears [35], the lack of any pulsatile signal in Figure 2.1 shows its minuscule size

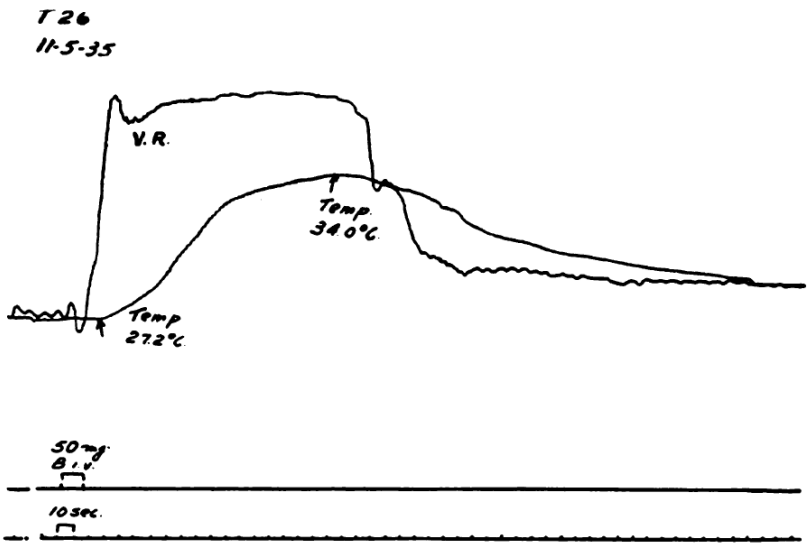

Figure 2.1: *Extract from Molitor and Kniazuk's article [34] (Figure 4, page 11). The small oscillations in the V.R. (vascular response) trace are not due to the PPG as the timebase would put their frequency around 0.1 Hz.*

compared with the vascular changes caused by the pharmaceuticals (if the pulsatile PPG's amplitude had been large enough, the kymograph input bandwidth was sufficient to display it). In the second half of the PPG trace, oscillations occur with a periodicity of approximately 0.1-0.15 Hz; these are likely to be Traube-Hering (Mayer[1]) [37] waves caused by internal periodic vaso-constriction and are known to often exceed the amplitude of the PPG [38, 39].

In addition to artefacts caused by the 'subject-under-test' (due to movement), external sources of artefacts were mentioned - albeit briefly - by Hanzlik et al. (1936): *"There should be sufficient protection of the photronic cell to prevent galvanometer deflections traceable to changes of illumination from outside sources."* [33]. At this point in history, there *was* sufficient technology and knowledge available to electronically counteract such disturbances as ambient light, but as the technique was used entirely for research purposes, there was no need as experiments could be carried out in controlled conditions.

[1]Multiple low-frequency waves occur in the vascular system, the slow arterial-pressure oscillations at 0.1-0.15 Hz as described originally by Meyer are not linked to the 0.3 Hz modulations described elsewhere in rabbits [36].

In both of these cases (subject movement and ambient light), the solution to the artefact problems was to remove the artefacts at the source such that there was no effect on the PPGs to begin with. For these particular experiments, this was not a problem as all variables that could cause such artefacts were controllable. However, as the use of photoplethysmography slowly progressed from experiments in laboratory conditions to more 'consumer-level' devices, the controllability of the sources of artefacts became more difficult. Now PPG devices are appearing in devices used for recreation, sport and high-motion environments; the artefacts are more prominent, and in response, the techniques required to remove or suppress such errors need to be increasingly sophisticated[2].

2.2 Applications

Recently, there has been a strong trend for commercial designers to incorporate technology into clothing and other wearable accessories; indeed, an entire industry in its own right has emerged [40, 41]. Photoplethysmography - particularly in-contact modes - are (in theory) ideally suited to being part of this trend as the continuous monitoring of a person's health is desirable. 'Smart watches' that contained the capability for data processing and transmission were ideal candidates for the inclusion of PPG sensors, helped by the fact that there was almost constant skin contact with the device. However, as with many commercially focussed designers, the technology was included in products without a sufficient understanding of the problems that can occur, leading to poor initial results...

In 2014, one reviewer of wrist-band/watch based PPG heart-rate monitors stated for the five new devices that were tested: *"To accomplish these readings, optical sensing requires you to hold absolutely still – no talking, no*

[2]Whilst it must be noted that these applications require much less fidelity in the measurements, artefacts re still a problem and must be removed.

moving, no muscle-tensing, no sweating, no smudging allowed. Thankfully, these devices will tell you when you're too active for them to work." [42]. This problem of motion-based artefacts was slow to be picked up on by designers. However, by mid-2015, products were being sold that performed with increased accuracy. *"If you are in the market for a running watch and heart rate monitoring is of secondary importance then these devices are fine; you will get a reasonable estimate of your average heart rate for a run"* [43]. The tests that the reviewer ran showed improved correlation between the wrist-band optical heart-rate monitors and the standard ECG (see Figure 2.2).

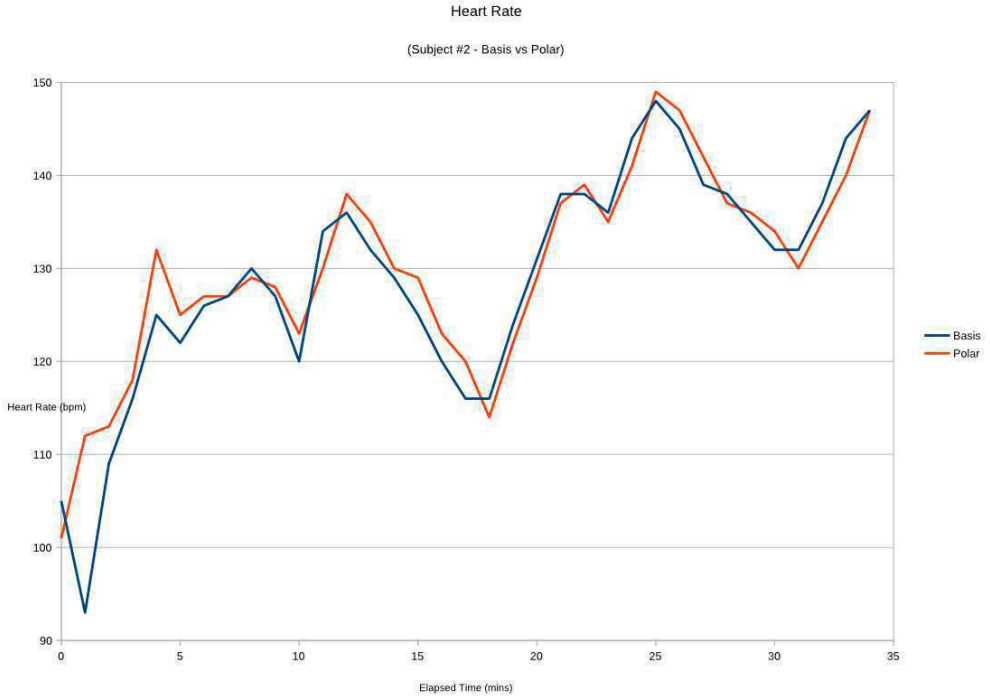

Figure 2.2: *Heart-rate trend comparison between standard ECG monitor (from "Polar") and a PPG wrist-band heart-rate sensor ("Basic Peak"). From online article [43].*

It was, however, also stated as a conclusion to the article that *"...if you are doing training based on heart rate zones or you want the most accurate data possible then we cannot recommend them"* [43], which shows that accuracy is still a concern. Additionally, although it is known that many devices do output heart-rates more frequently, one of the devices that

was tested appeared to only output a heart-rate every minute (Figure 2.2). For sports and fitness applications this is often more than sufficient as it is the long term trends that are more important.

Many medical applications require additional information to just the heart-rate. For those where the heart-rate is important, a value output more than once per minute is necessary; additionally, any significant delay would be considered unsatisfactory [44]. The shape and amplitude of the PPG, and the position in time of individual pulses, provide the most information that can be gained from the technique (using a single channel). When no motion is present, such a signal is trivial to extract; fewer than six fundamental electronic components and a data logger or oscilloscope are all that is required (Figure 2.3).

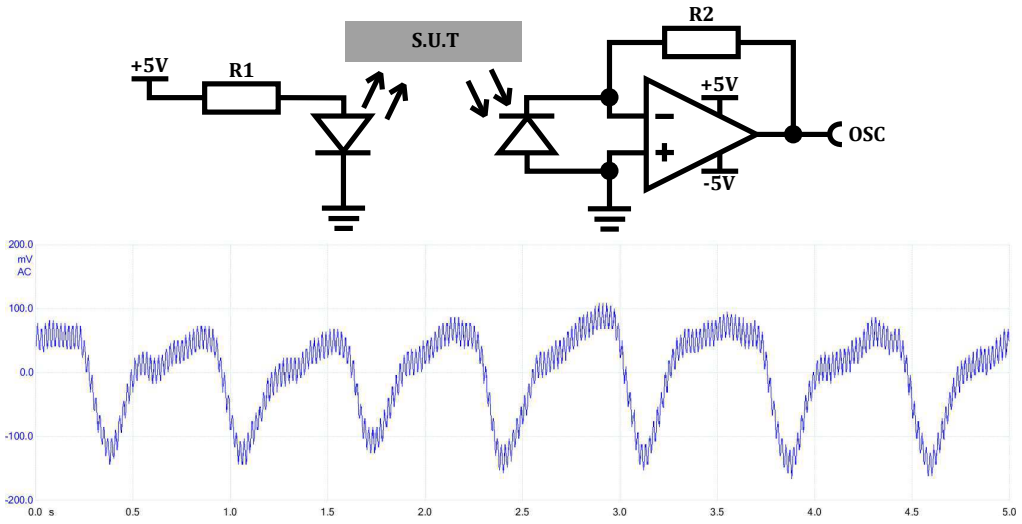

Figure 2.3: *Example of a simple - trivial - circuit required to extract a PPG signal, including the resulting signal (Oscilloscope is AC coupled to remove 'DC offset'). Note that 50 Hz 'mains hum' exists in the output as no filtering is used.*

When motion is present, however, the problem becomes far from trivial. Many attempts have been made to suppress the effects of motion artefacts from photoplethysmographic measurements (see Section 2.4), and many have succeeded for their particular application. However, the techniques that have been developed do not appear to be universal solutions as slight adjustments to the application set-up often causes much degraded results.

Although the primary motivation for this research came from a specific application (neonatal resuscitation), the goal was to create a solution that could be applied to multiple fields with similar success.

2.2.1 Neonatal Resuscitation

From previously collected statistics, both from within the UK and internationally, it has been shown that approximately one in ten newborn babies require resuscitation directly after birth in one form or another [45, 46]. This resuscitation is performed manually by trained personnel (midwives/-clinicians) after a problem has been detected, and the procedures, although fairly straightforward, require experience and considerable skill to apply [47]. Of all the signs that the person performing the resuscitation is required to keep track of, it has been quoted multiple times that the heart-rate takes the highest priority as it allows for the fastest indication of recovery, for example, *"Heart Rate - This was found to be the most important diagnostic and prognostic of the five signs."* by [48], and *"An initial assessment of heart rate is vital because an increase in the heart rate will be the first sign of success during resuscitation."* by [49].

Through work done at the University of Nottingham, combined with prior knowledge at the time, a company was created to develop a hardware device to allow for fast, reliable and accurate measurement of the heart-rate of newborn babies [50, 51]. This company, SurePulse Medical Ltd.[3], co-founded by the University of Nottingham and a Derby (UK) based company *Tioga Ltd.*, originally based the design off of a previous student's work at the University, before developing it further into a saleable product. This product aimed to use the current 'state-of-the-art' for the PPG extraction technique.

One example of previous work at the University of Nottingham regard-

[3]Formerly HeartLight Systems Ltd.

ing photoplethysmography in medical settings, mainly in the resuscitation of newborns directly after birth, was looking specifically at optimising the wavelengths used for PPG. It is common for pulse-oximeters (that optically measure the oxygen content of the blood) to use two wavelengths of typically red and infra-red [52, 53], and the devices that also output a PPG signal would often use either/both wavelengths to facilitate this. However, this is not optimal as the PPG signal is heavily dependant on the wavelength used [54]; green (540 nm in reflection-mode, 510 nm in transmission) was found to be the optimal wavelength to use.

Other work at the University included looking at improving analogue front-ends, the design of the sensor head itself, and the methods used to extract additional information from the PPG - including the blood-oxygen saturation (this also included the development of ultra-miniaturised sensors on silicon integrated-circuits [55]). However, the problem that was presented at the start of *this* research had not been looked at in detail before (in the University) and offered new challenges as well as many opportunities and potential to be included in the SurePulse device being developed: motion artefacts. As motion artefacts have been (and still are) a considerable problem for most PPG measurement applications, a novel method of removing such effects was the goal. Although the original motivation came from the aforementioned medical application, the hope was that a universal solution could be created.

2.3 Comparison with ECG

Multiple methods of measuring the heart-rate are in use today. For single one-off measurements (that are taken once to get a general idea of the health of a person), medical staff may decide to use a stethoscope and manually count the audible heart beats in a given time window, or do

the same with palpations in the wrist or neck [4]. This method is - of course - prone to human error; however, for times when high accuracy is not required, this is acceptably cheap, simple and fast. When long-term or continuous measurements are required, the manual aspect of using a stethoscope becomes overburdening and unacceptable: an automated method is required.

The *de-facto* standard currently used for continuous measurement of heart-rate is the electrocardiogram (ECG) [56]. In hospitals, this typically takes the form of a unit connected to the patient via at least three electrodes. If additional information about the patient is required other than the heart-rate, such as morphology, then more than three electrodes may be used (a '12-lead' ECG is commonly used[4]) [57].

2.3.1 Advantages Exclusive to PPG

Advantages of a PPG based system are almost entirely mirrored by ECG's disadvantages, and are part of the reason why research into using PPG is currently so strong. Photoplethysmography requires no galvanic connection between the patient and base-unit; it is therefore inherently safe and trivial to get high isolation voltages. Additionally, each sensor is a self-contained unit requiring just a single cable; even if transmission modes are used, the light sources and detectors are mounted in the same module: typically a finger-clip [58]. The possibility of performing non-contact (remote) PPG is also advantageous as this allows for multiple detections sites, either on one or multiple people simultaneously when imaging is used.

Even though it is possible to perform 'non-contact' electrocardiography [59–61], the technology is notoriously sensitive to electrical noise and motion artefacts, limiting its use to a few potential applications. Unless

[4]Despite the name, a 12-lead ECG only has 10 physical connections (RA, LA, RL, LL, V1-6). By looking at the differences between combinations of the electrodes, 12 individual measurements are made.

the sensitivity of the systems are increased it is unlikely to be used for healthcare monitoring of newborns, for example, as the electrical signal is so weak. Remote or 'non-contact' photoplethysmography (there is no distinction made in this thesis as there is elsewhere) is not only possible, however, but fairly easy to accomplish when no motion is involved [62–64].

Due to the mechanical simplicity of photoplethysmography, the time taken to set up the necessary equipment and connect it to the patient is far less than for ECG. This allows for the heart-rate (and other vital signs) to be measured and displayed more quickly, allowing for higher chance of survival in time-pressured situations [44].

SpO_2 (blood oxygen saturation) estimation can also be performed with the appropriate pulse oximetry equipment, without the need for in-vivo probes [52].

2.3.2 Disadvantages of PPG

One of the limiting factors with PPG is that the power consumption of the device will be higher than ECG. This is due to the power required to illuminate the light sources used in the measurement. ECG is a passive[5] sensing method that does not require any more current than that of the controlling and sensing circuitry. Techniques exist that can reduce the power consumption of the PPG system (by pulsing the LEDs, for example), but the battery lives of portable devices will still suffer.

Locations from which PPGs can be extracted from a human body depend on the mode of operation of the sensor. As described in Section 1.2.3, transmission-mode set-ups require a short optical path between opposing sides of the body (fingers, toes, ear-lobes, etc.) and therefore have limited applicable sites to be placed. Relative comfort between the two methods

[5]ECG uses an output lead typically called the 'right leg drive' to set the patient's voltage to a function of the inputs. This is used to make a differential measurement with increased noise suppression (the noise is cancelled out). This drive output is a voltage bias and uses/sources negligible current.

will be subjective and heavily dependent on the patient. That said, the increased number of leads and electrodes required for ECG will negatively contribute to the comfort. On the other hand the PPG finger clip will potentially restrict movement and the use of the hand.

The focus of this thesis addresses one particular disadvantage of PPG: motion artefacts. PPG is particularly prone to errors caused by mechanical movement as this disruption alters the optical path between the source and sensor (specifically, for in-contact modes, the non-homogeneous skin and tissue is displaced or compressed, causing variations in absorption).

2.3.3 ECG/PPG Commonality

Despite their differences, there are many applications where PPG could replace ECG without any loss of information or validity. For medical applications specifically, the feasibility of using PPG instead of ECG has been looked at in detail by many researchers world-wide...

One particular study looked at the detection of 'hypertension' (high blood pressure) using the two methods [65]. The results showed that whilst a deviation in blood pressure could be determined by looking for changes in the ECG-pulse to PPG-pulse time, known as the pulse wave transit time, could be found, either the ECG or PPG pulse intervals could be used as an indicator of hypertension. This study showed that when the pulsatile signal frequency (and occurrence) is required for a particular technique, information extracted from the PPG is sufficient.

Another study looked at the pulse occurrence (frequency/period) to determine another metric that is commonly used within clinical settings: heart-rate-variability (HRV). Again, as this technique does not require information that is unique to either PPG or ECG, the results were satisfactorily promising. The authors concluded with *"...PPG signal could be as dependable as the popular ECG in the derivation of the HRV signal."* [66].

Research by Lu et al. (2009) has also looked at heart-rate variability with the same results: *"...PPG recording could provide more clinically valuable information..."*, *"...Our results support this proposition by demonstrating that HRV analysis of signals derived from 3-lead ECG and earlobe PPG recordings are almost identical"* [67]. However, in their closing statement, one interesting point was raised: *"The potential advantages of PPG over ECG to derive HRV warrant further investigation in both ambulatory patients and those whose treatment is likely to generate electrical artefacts in their electrocardiogram."* Some treatments may require forms of continuous electrical stimulation (i.e. post-surgical stimulation of muscles [68], or de-fibrillation), in which case an ECG signal may easily become corrupted. In the case of post-surgical stimulation, the patient will be in a vulnerable state having recently had an invasive procedure; it would be reasonable to think that the heart-rate would be closely monitored during this time to watch for signs of improvement or deterioration. The originally suggested problem that some treatments would generate continuous electrical artefacts could be easily solved by using the entirely optical technique, PPG. For short duration electrical interference generated by, for example, de-fibrillation, ECG signals are often 'blanked' to remove the artefacts.

In summary, some treatments or monitoring techniques require the use of either ECG or PPG exclusively due to the fundamental nature of the source of the information; electro-cardiology will only use ECG as it is the heart itself - and particularly its electrical activity - that is being monitored. Monitoring of blood perfusion and circulation, on the other hand, will fundamentally use optical techniques such as PPG as there is no electrical connection with the heart. However, it is clear that many ECG-based techniques that are currently in use could be equally effective using PPG - particularly those where the pulse occurrence and frequency (heart-rate) are required. The advantages of PPG allow for increased convenience,

simpler designs and immunity to electrical interference. The primary disadvantage to overcome is the susceptibility to motion. Due to the benefits of PPG, many researchers have attempted to solve this problem of motion artefacts, with varying success.

2.4 Existing Research

Analysing techniques and results from other researchers revealed several important factors that were of use when designing a novel solution to overcome motion artefacts. These factors were not necessarily *directly* related to motion artefact reduction, as many contributing factors can either help or impede the development of such a system. Many previous works involved looking at widely varying aspects of PPG extraction, although some common themes were observed. These include: extensively using laboratory conditions and 'ideal' set-ups to show better results, ignoring how the techniques would work with real-world data; a predisposition to using particular filtering techniques; and using methods that - although recover a great signal - will have little practical application due to prior requirements of the technique. In fact, it was discovered that some published work appeared to entirely neglect some aspects of prior knowledge in order to present more impressive results. A selection of studies are presented below that helped to identify which methods and techniques were beneficial to study further.

2.4.1 Frequency Components

When two independent signals that occupy different frequency bands are present in a recording, it is fairly straight-forward to filter either of them

out with a simple filter[6]. Care has to be taken with designing the filters, of course, but as long as the filter roll-off is sufficient, such artefacts (with different frequency components) can be removed.

When a published paper states that when collecting their data: *"The frequency of induced artefact was controlled to be at least 3 Hz, and data sets that did not meet this criterion were discarded"* [14], the technique that was created/used has no guarantee of working with data-sets that were of the opposite criterion. Figure 2.4 shows a figure from the paper illustrating the frequency components of one measured signal.

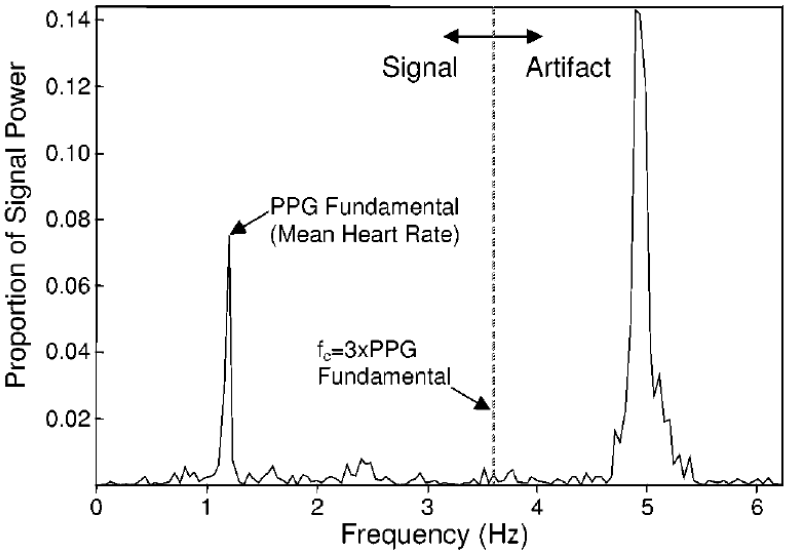

Figure 2.4: *Extract from Hayes and Smith's paper (2001) showing the enforced frequency separation between the PPG and artefact fundamentals ([14], Figure 4).*

The reason given for the choice of induced artefact, and specifically its frequency, was that it was their requirement to quantitatively track the two components (the PPG and the artefact) throughout the use of the developed algorithms. As neither the PPG nor the artefact were particularly controllable, they could not be easily characterised; therefore, if they occupied the same frequency range, there would be no way of knowing how

[6]It is understood that if the band-widths of the two components are not zero, some overlap may occur; however, with simple cyclic signals, the majority of the amplitude is contained within the fundamental. Therefore such an approximation is accurate.

effective their technique was. This reasoning is understandable as the non-uniform nature of 'real' artefacts makes characterisation difficult. However, the problem of removing artefacts from PPG signals cannot be implied to be solved due to the result as there is a significant possibility that the technique will not work for signals where the PPG and artefact components overlap in the frequency domain.

The authors did acknowledge this, and suggested that their technique would in fact **not** function correctly given the previously mentioned 'real-life' conditions: *"Removal of PPG artefact cannot be easily performed by signal processing methods because of the likely spectral and temporal similarity between signal and artefact, which may be compounded by statistical interdependence or extreme differences in the magnitudes of signal components due to artefact and the arterial pulsations"* [14].

Hayes and Smith's technique is one based around using additional wavelengths to provide information about the motion as seen from the photo-detector. As the study is focussed on improving pulse-oximeters, the original two-wavelength system (red and infra-red) is replaced with a three-wavelength system which introduces a new light source at around 800 nm where the absorption due to haemoglobin and oxyhaemoglobin is approximately equal (meaning a resulting signal that should be independent of the blood-oxygen saturation). Ratios of each of the red and infra-red signals to this new wavelength's are used in the algorithm. In the paper's discussion section, the types of motion that the technique works best with are mentioned along with an explanation; this motion is the result of the application of pressure to the sensor. This results in *"modulation of the received intensity"* which can be countered by the developed model easily. Problems show themselves when significant distortion of the area under the sensor occurs, particularly when there is *"redistribution of venous blood"*, as they state.

In summary, out of 'motion magnitude' and 'frequency selection', the only combination that Hayes and Smith can show their technique solves is one where there is a small amount of motion, and where the frequency of the artefact is outside the PPG's bandwidth (three harmonics of, as they state). For more vigorous motion, their results still show some improvement, but the remaining artefact will still frequently corrupt the PPG to an unacceptable level. For the research presented in this thesis, it was decided that direct characterisation of the artefacts generated would not yield meaningful results; instead, using the PPG with no motion as a baseline, the deviation from this baseline (due to artefacts) would be a sufficient measure of corruption. Because of this decision, the more realistic scenario of having artefacts overlapping the PPG in the frequency domain will be used as this is the condition that is most likely to occur in the target applications.

2.4.2 Application Conditions

Motion artefacts are more prone to occur in non-contact variations of PPG than their in-contact counterparts; this is due to the lack of mechanical coupling between the sensor and the subject allowing for greater freedom of movement. For non-contact PPG, certain types of motion will have more effect on the signal than others; the most significant corruption will be caused by a panning motion as a result of the topology of the skin [69]. When the motion takes the form of a rotation (the detection point on the subject being the rotation centre), then the artefacts are sufficiently *well-behaved* to be able to suppress.

Cennini et al. (2010) published a paper looking at countering motion artefacts caused by a rotation of a subject's hand [26]. This PPG measurement technique used multiple wavelengths in a similar fashion to Hayes and Smith's paper previously discussed. Despite this non-contact (by around

30 cm) method using two wavelengths, as it was the PPG itself that was of interest, no pulse oximetry was planned. An assumption the authors made was that the signal they received containing the PPG was linearly mixed with the artefact; the adaptive echo cancellation method they used to remove the measured artefact is most suitable for this type of mixing.

The fundamental principle behind the proposed technique is based on the fact that different wavelengths will be absorbed dissimilarly by different parts of the skin; the shorter, blue, wavelength used (450 nm) will be absorbed more readily by the skin's surface, compared to the infra-red (970 nm) which will scatter beneath considerably more. In theory, the signal obtained by the longer wavelength will contain both the PPG and any present artefact, whilst the shorter wavelength should mostly contain artefacts[7]. This technique makes the assumption that the PPG signal itself is intact behind the artefact, and that by subtracting an artefact 'reconstruction', the PPG can be recovered. This is what is implied by linear mixing. In other words, the proposed system's outputs have no redundancy in them; the PPG is only present in a single channel, so if that becomes corrupt as an *indirect* result of the artefact, it is impossible to recover.

As a result of this, it is believed that if the motion were to be changed from rotation, where the region of interest (ROI) on the subject is stationary, to 'panning', where it moves, then the PPG signal itself will be corrupt independently of the artefact source - this is due to the two wavelengths no longer pertaining two linear mixes of the PPG and artefact.

The main point to extract from this paper is that for some applications (and more particularly for some motion types), systems exist that are able to suppress artefacts with relative effectiveness. However, the same technique has no guarantee of working with other types of motion. Linked to

[7]Practically, some light will be scattered back to the same side as the point of entry, which will be detected as the pulsatile signal. This is visible in their paper ([26], Figure 6).

this, using extra wavelengths to estimate the motion *as seen by the photo-detector* may not always provide a means of compensating for motion.

2.4.3 Adaptive Filters & Accelerometers

As motion artefacts are the largest problem faced when extracting pulsatile signals using light, it would be reasonable to think that the motion could be compensated for using a direct measurement of the motion itself. Many researchers have looked at this proposition and have used a similar approach in their techniques [70–72]. Accelerometers are used to measure the motion close to the PPG measurement site, and then the PPG signal is filtered using an adaptive filter that is 'trained' in real-time by the accelerometer.

These adaptive filter techniques all require some form of 'cost function' which is used to optimise the filters to achieve the best filter configuration. For the purposes of rejecting artefacts from a PPG signal, this cost function attempts to maximise the similarity between the measured signal and that obtained from the accelerometer (by filtering it). In other words, the initial goal is to filter out everything *but* the artefact. The resulting signal which now ideally should contain no PPG can be subtracted from the original signal; thus, the true PPG should be revealed.

Lee (2010) states that: *"the transfer function from motion to artefacts in a physiological system cannot be easily described by a definite physical model"* [72]. This is one of the reasons why a fixed filter is not suitable for PPG artefact rejection. However, this statement also alludes to a problem with adaptive filters. As, at their core, they are standard temporal filters, the PPG signal can still be erroneously filtered out if the PPG and artefact frequencies collide, or if the signal used to train the filter is not an accurate representation of the artefact. As the physical motion (as measured by the accelerometer) and the resulting artefact (as seen by the photo-detector) are not exactly the same, the filter which is 'trained' using the former will

never be able to completely remove the latter. This technique can therefore only work with some types of artefact that have a very strong link between the actual and observed motion.

In summary, the relationship between the physical motion that causes the artefact and the artefact itself is - obviously - related, however, the function that converts one to the other is highly complex and perhaps even indeterminable. Accelerometers can therefore inform about when motion occurs in order to apply further processing, and in some circumstances be used to train applicable filters, but must not be used to directly modify the measured signal. Relating to this, adaptive filters may be used to filter artefacts, given enough information about the erroneous signal, but as the filters are basic in nature, they will not distinguish between artefact and PPG if they both occupy the same frequency range, as described in Section 2.4.1.

2.4.4 Processing & Unsuitable Filters

Wu et al. (2012) from MIT, along with techniques to effectively amplify motion, presented a method for amplifying small changes in colour. The technique looks for small variations over large areas using spatial decomposition, before re-applying (summing) the results locally on the data. This allows for *sub-quantisation level* measurements; i.e. the detection of changes smaller than the 'least-significant-bit' of the sensors/ADCs used. One of their claims is to be able to visualise, and therefore extract, the cardiac-synchronous signal from a video in a similar fashion to other researchers. However, with the colour amplification technique, this becomes much more effective.

The source code written in MathWorks^{TM} MATLAB is provided by the researchers (linked from within their article [73]). After inspection, it is clear that despite the original claim in the accompanying article that

"we automatically select, and then amplify, a band of temporal frequencies that includes plausible human heart rates" [73], this isn't what actually happens. Instead, a very narrow (0.17 Hz), and ideal *top-hat* band-pass filter is applied to the data before it is processed further.

In a later part of the article, the 'true' objective of the research becomes clear when they state: *"We first select the temporal bandpass filter to pull out the motions or signals that we wish to be amplified."* [73]. The technique the researchers at MIT have developed is not one of movement and colour *extraction*, but more one of *visualisation*. Although not tested, it is believed that any motion that exists in the video will corrupt the amplified signal; this is because it is the colour variations that are being used to amplify the colour and motion will cause a similar colour variation. It is exactly this process that they use to amplify motion.

Although the original claim failed to be realised in the provided code, one very important point can be extracted from this research: filtering of raw signals. If the desired signal's frequency is known[8], it would be too easy to simply band-pass filter the signal to remove everything else. If, in the more likely scenario, the desired signal's frequency is not exactly known, then a sufficiently wide band-width is required to encompass the expected range. For the heart-rate fundamental frequency and the harmonics thereof, this bandwidth is surprisingly high. Section 2.6.2 looks at the problems of filtering PPG signals more closely; it was the results in Wu et al. 's paper [73] that prompted this closer look.

2.4.4.1 Filter Bandwidths used in Literature

Band-pass filters, used to reject frequencies outside of a selected range, are commonly used in PPG measurement circuitry and digital post-processing to enhance the signal-to-noise-ratio (SNR) and/or signal-to-artefact-ratio

[8]Then a technique to extract the heart-rate is not even needed.

(SAR) before further processing. For many applications the filter design is critical to ensure that the PPG is not distorted. One critical factor when designing such a filter is the bandwidth. Table 2.1 lists general filter specifications from several published papers and documents.

Reference	**Range**	**Band-width**	**Type**
Lee et al. (2010)	0.50 - 10.00 Hz	9.50 Hz	Butterworth[1]
Unknown (2006)	0.30 - 8.00 Hz	7.70 Hz	Butterworth
Bahga et al. (2010)	0.50 - 5.00 Hz	4.50 Hz	Moving Average
Nguyen et al. (2011)	0.50 - 5.00 Hz	4.50 Hz	Butterworth
Grubb (2009)	0.50 - 5.00 Hz	4.50 Hz	Unspecified
Wu et al. (2012)[2]	0.40 - 4.00 Hz	3.60 Hz	Ideal (Top-Hat)
Zhang (2015)	0.40 - 4.00 Hz	3.60 Hz	Butterworth
Cennini et al. (2010)	0.75 - 3.50 Hz	2.75 Hz	Unspecified
Verkruysse et al. (2008)	0.80 - 2.00 Hz	1.20 Hz	Butterworth
Wu et al. (2012)[3]	0.83 - 1.00 Hz	0.17 Hz	Ideal (Top-Hat)

Table 2.1: *Table of filters characteristics used in literature to first 'improve' PPG signal, sorted descending by Bandwidth.*
([1] Assumed through description, [2] Described in text, [3] Used in practice).

With the assumption that a typical heart-rate range (for all age groups) is between 45 and 210 bpm, and that up to three times these frequencies contain significant harmonics of the PPG signal, the appropriate initial filter should be between approximately 0.75 Hz and 12 Hz. The actual bandwidth of the PPG will be higher than this, but the higher frequency components (harmonics) will have considerably smaller amplitudes, and can be ignored. The use of filters in PPG acquisition and processing is discussed later (Section 2.6.2).

Some of the PPG waveforms presented in the documents tabulated appear highly distorted. As explained previously, this may be acceptable for heart-rate extraction, but if further processing is performed to extract more subtle variations, the resulting information will be meaningless. That said, with some types of filters (such as Chebychev etc.), the pass-band will not be 'flat'; when heart-rate detection algorithms are used, this may result in biased outputs as some frequencies will be attenuated more than others.

2.4.5 Existing Research Summary

After careful observation of the work of other researchers, several key points emerged that are of key importance to the research undertaken in this project.

Real-World Data Using generated artefacts and/or PPG signals will be useful in some situations where algorithms will be tested or where specific applications are used. However, as the research presented in this thesis is looking at a novel method of suppressing artefacts, no specific application will be targeted (despite the reason for this project existing being a specific application), and so real-world data must be considered.

Filter Selection Many researches use filters that have far too narrow pass-bands to extract the 'correct' PPG. Some have reasonable excuses for this, whether it being heart-rate (fundamental) extraction or for visualisation. However, it is acknowledged widely that an incorrect filter (particularly one that is too narrow) can corrupt the PPG's shape, or indeed remove vital features that would be useful/necessary in extracting more subtle information, or indeed reducing artefacts.

Existing Techniques Limitations Attempts at reducing artefacts using alternate external sources of information, such as with accelerometers or additional wavelengths, have limitations. For accelerometers specifically, the actual source of the artefact will present itself differently between two different acquisition methods: the accelerometer and the optical photo-detector. The actual sources of the motion must be considered when being used in correction techniques.

These points were taken into account in the following research as these were considered the main faults (or sources of error) in existing research.

2.5 Hardware

The hardware portion of a PPG acquisition system may be the most important part as even with the best processing available, a poorly acquired signal (one with large amounts of noise, or where the data clips, for example), will not be able to yield the desired result.

2.5.1 Wavelength

The wavelength of the light sources used is an extremely important factor when detecting the PPG and using it to make physiological measurements. For pulse oximetry, the blood-oxygen-saturation (SpO_2) can be calculated from variations between signals measured using different wavelengths; the different absorption curves of haemoglobin and oxyhaemoglobin are what facilitates this (see Figure 1.5).

For heart-rate extraction, which is the goal of this project, a single wavelength is the *minimum* that is required, and the optimum wavelength is the one that provides the strongest signal-to-artefact ratio. There exists no known system that is completely insensitive to motion (for photoplethysmography), so the best solution in this case is to maximise the ratio: the maximum response to changes in blood volume for the minimum change caused by the possible variations (caused by motion) in the surrounding tissue. Although the actual reasons for such behaviour is complex and not trivially modelled, published empirical findings exist that show that wavelengths around the 'green' range ($\approx 550\,\text{nm}$) provide the best results [54, 74] (this is referring to the absorption of light by the blood).

It is this wavelength, or one close to it, that it predominatly used for reflection-mode PPG in research [75]. By chosing a different wavelength, the opportunity for poorer quality signals may become in insummountable problem when processing the recorded data.

2.5.2 Sensitivity

Pulsatile blood is not the only absorber in the optical path between the light source and sensor. As explained in the Introduction, light is absorbed by melanin in the skin, as well as muscle and the surrounding tissues. The compression of these other elements as a result of motion creates a change in absorption (or scattering), resulting in artefacts. The amount of light that falls upon the sensor will vary heavily depending on the location of the measurement and on the subject themselves.

For subjects with darker skin, due to an increased volume of melanin, any light that falls upon the skin will have a higher chance of being absorbed. As this absorption is constant, this will result in a reduced detected light intensity. However, as the pulsatile signal is the result of the absorption due to blood, different skin colours will only affect the ability to extract a PPG signal due to a reduced light level [74]. Additionally, some locations on the body will have better blood perfusion than others, resulting in variations in both the 'AC' and the 'DC' components of the PPG. Finger-tips, toes and ear-lobes are locations with high perfusion (also, conveniently, peripheral to the body making easy access for transmission mode PPG). Locations where dense capillary beds exist provide the largest AC components [76].

For either case, to accommodate for the different average light levels, the sensitivity of the sensor must be sufficient to detect the smallest AC components within the largest DC components; in other words, a high dynamic range. When weak PPG (AC component) signals are encountered, properties of the photo-diode will limit the ability to extract the waveform. Conversely, when the DC levels are high - often caused by ambient light falling on the sensor - it is the electronics and conversion circuitry that limit the ability to obtain the PPG.

2.5.3 Brightness

To overcome the issues regarding sensitivity of the photo-detector in order to detect smaller changes in the absorbed light, it would seem reasonable to simply increase the amount of light that enters the system. This, however, causes two main problems.

Firstly, the pulsatile and non-pulsatile components will both be increased proportionally with the input brightness; for the pulsatile PPG signal, this is an advantage and the desired effect; however, for the non-pulsatile 'DC' component, this will more likely saturate the electronic amplification in the following stages. Any attempt to remove the non-pulsatile (DC) signal in the optical path (via attenuation) will also attenuate the PPG (i.e: you cannot AC-couple light). It is clear that this problem must be solved in the electronics conversion stage.

Secondly, as the system is illuminating the skin, potentially at close range, safety has to be considered. As the light sensor has the opportunity to become detached from the skin for in-contact modes, and is 'free-floating' in non-contact modes, the user could look (intentionally or accidentally) directly at the sources; this has the potential to cause eye damage. Some wavelengths of light are more harmful than others, but even 'safe' wavelengths can be dangerous at high intensities. Regulation IEC/EN-62471 covers the acceptable powers and exposure times for different wavelengths (and is summarised in [77, 78]). For the green wavelengths primarily used in this research at around 550 nm, there are three types of hazard: *retinal blue-light*, *retinal thermal* and *thermal skin*. However, due to the very low power levels used in this research, the LED sources are exempt, with an unlimited exposure time.

2.5.4 AC or DC Coupling

To simplify the electronics and conversion stages, some researchers choose to 'AC couple' the signal from the sensor [71, 72]; this has the effect of removing the DC and low frequency components from the signal, and has the advantage of preventing the following amplifier stages from saturating due to large DC components. A PPG signal will have some information in these lower frequency components, mainly due to breathing and the aforementioned Meyer waves, but the pulsatile signal from which the heart-rate can be extracted will not be lost. Counter-intuitively, perhaps the largest problem with this technique is that potentially vital information about the *artefacts* will be lost, preventing efficient heart-rate extraction later.

'DC' coupled systems (for example, as used by [34, 79]) have to be designed more carefully to prevent saturation of the amplifiers and converters, but as the true signal is passed through to the signal processing stage in its entirety, instead of a filtered one, there is the potential for better filtering techniques to be used (ones that may not be possible to implement using traditional analogue electronics, for example), and therefore allow for improved results. Cameras use DC-coupled systems as their primary purpose is to capture the absolute intensity of light. Due to their architecture, cameras also digitise the input *on-chip* such that the output is already digital.

2.6 Data Processing

Even if the electrical and mechanical stages of the system were perfect in all respects (perfect mechanical coupling of the sensor and skin, and an 'ideal' ADC with no pre-filtering required), there would still be a need for signal processing of the data. This is due to the fact that the detected light

will contain additional components alongside the desired PPG. These components can broadly be classified into two groups: 'noise' and 'artefacts'. The definitions of these two can vary considerably between fields and even between authors. The definitions used in this thesis are as follows:

Artefact A *signal* that originates from outside the measurement system, that is not desired.

Noise Random and unpredictable fluctuations within a signal that are of no practical use to the system.

In all cases, noise is both an undesirable component of the data, and one that cannot be used to extract the PPG. Noise is random in nature and generally has a broad frequency spectrum. This makes it impossible to independently measure (for the purposes of subtraction from the signal). Artefacts on the other hand, despite also being an undesirable component, are signals in their own right with defined 'sources'. In theory, if the 'artefact' became the desired signal, then the component that was originally the goal for extraction would then become the artefact. If the artefact *as seen by the photo-detector* could be extracted accurately, simply subtracting it from the originally obtained signal would reveal the PPG. Unfortunately, this is rarely the case.

2.6.1 Modulation

By turning the light source on and off at a known frequency, it can be detected more easily when in the presence of external signals and noise; this is one of many techniques that falls under the term 'modulation'. For this example, specifically, the process is known as amplitude modulation (AM), where the amplitude of the 'carrier' (the known frequency signal, in this case) is perturbed by the source signal. For PPG, the light amplitude is directly modulated (attenuated) by the pulsatile blood in the skin.

When two frequency components are multiplicatively modulated (mixed) together, new frequency components are formed as the sum and difference of the two. With suitable filtering and choice of carrier frequencies, this effectively allows for the desired signal to be moved away from other nuisance signals (mains hum, optical flickering from mains lighting and computer monitors, etc.) and noise ('one over f'/flicker noise [80], etc).

When the photo-detector detects the modulated light (the product of the light source carrier and the attenuation due to the PPG), additional light that is not modulated from other sources may also fall upon the sensor. This externally sourced light may, or may not, be attenuated by the same means as the controlled light source; regardless, it is *added*[9] to the modulated signal. If no modulation was used then this external 'artefact' would corrupt the desired signal.

[9]This is assuming that the detector has a linear response to light intensity. Non-linearities in the detector could result in the signal components mixing.

2.6.1.1 Example

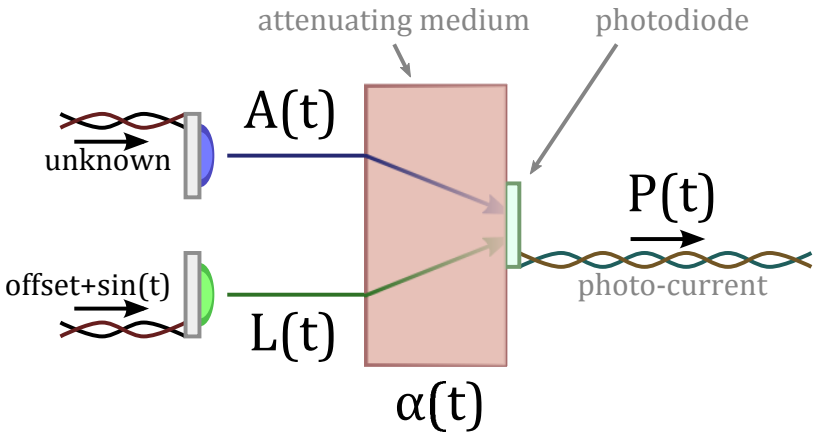

Figure 2.5: *Simple optical system used to demonstrate the advantages of modulation.*

An example optical system is presented in Figure 2.5 with a controlled light source (an LED with intensity $L(t)$), an **external** uncontrolled light source (the artefact, with light intensity $A(t)$), and a time-varying attenuating medium (representing a finger with pulsatile blood-flow, $\alpha(t)$), the optical output detected by a photodiode ($P(t)$) can be written as:

$$P(t) = \alpha(t)[L(t) + A(t)] \tag{2.1}$$

With a constant 'DC' light source (with intensity L_{DC}) and an artefact ($A(t)$), it is clear that the light detected by the photodiode would contain a combination of the two signals:

$$\begin{aligned} P(t) &= \alpha(t)[\cancel{L(t)}\, L_{DC} + A(t)] \\ &= \alpha(t)L_{DC} + \alpha(t)A(t) \end{aligned} \tag{2.2}$$

As the artefact function is not known, there is the possibility that it could either look like a PPG (like α), in which case it may erroneously be considered a PPG by a human or a machine; and/or it could corrupt the real PPG to such an extent that the PPG becomes unrecognisable.

With a sinusoidally time-varying controlled light source ($L(t) = L_{DC} + L_{AC}sin(2\pi F_m t)$), the output (the light detected by the photodiode) now contains an extra term (in **bold**):

$$\begin{aligned} P(t) &= \alpha(t)[L_{DC} + L_{AC}sin(2\pi F_m t) + A(t)] \\ &= \alpha(t)L_{DC} + \boldsymbol{\alpha(t)L_{AC}sin(2\pi F_m t)} + \alpha(t)A(t) \end{aligned} \tag{2.3}$$

Although the result contains the same components as in the unmodulated example, the new term represents a PPG-modulated sinusoidal carrier. This carrier is (when F_m is chosen as such) at a much higher frequency than either the PPG or the artefact, and a band-pass filter can be easily placed around it to remove the other components. To extract α from the remaining component, the signal must be demodulated.

Multiple techniques exist that can be used to demodulate/recover the signal. The simplest is to use an envelope detector, a technique that early AM radios used that consists of a lossy rectifier (Figure 2.6). The envelope of the AM signal is - by definition - the amplitude of the original signal. Alternatively, by simply rectifying the signal, and then performing a low-pass filter, the same signal can be recovered.

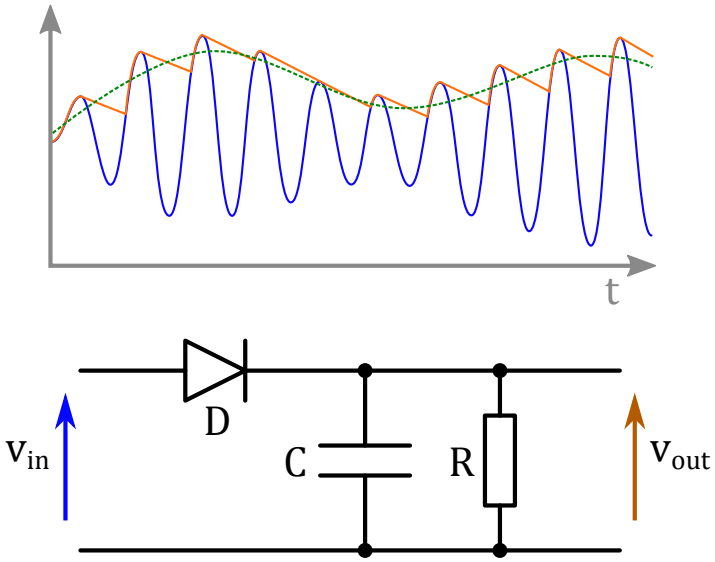

Figure 2.6: *Simple AM demodulator. The green dashed line shows the ideal demodulated signal.*

These examples assume that the artefact is an additive signal that is neither related to the PPG nor the controlled LED source. In reality, such artefacts occur externally to the PPG acquisition system and take the form of external light sources: office lighting, computer monitors, the sun, etc. For some of these sources that have predictable intensity patterns (the flickering of computer monitors and mains lighting, for example), a notch filter may be used to suppress them. However, in all cases, movement of the light source, or indeed the subject, will cause intensity variations in these signals that will have frequency components that fall within the PPG's range: filtering will not be effective. Modulation, or some form of external source correction, will correct for these artefacts successfully.

2.6.1.2 Ambient Light Subtraction

A simplified technique for removing the effects of external light sources, that is based of traditional modulation methods, is known as 'ambient light subtraction' (a practical form of 'Correlated Double Sampling'). In this technique, the controlled light source is turned on and off repeatedly at a relative high frequency (compared to the artefacts that may be present). The photodiode output is sampled twice per cycle, once when the light is off, where it just detects the effects of any ambient light in the vicinity; and once when the light is on, where it detects the same light, but with the addition of the effects of the controlled light source.

$$\begin{aligned} P_{ON} &= \alpha(t)[A(t) + L_{ON}] \\ &= \alpha(t)A(t) + \alpha(t)L_{ON} \\ P_{OFF} &= \alpha(t)[A(t) + \cancel{L_{ON}}] \\ &= \alpha(t)A(t) \\ P_{ON} - P_{OFF} &= \alpha(t)A(t) + \alpha(t)L_{ON} - \alpha(t)A(t) \\ &= \alpha(t)L_{ON} \end{aligned} \tag{2.4}$$

By taking the difference between the two values, the ambient light is effectively removed. This technique is much simpler to implement than the previously mentioned modulation methods, and indeed the necessary calculations in 'post-processing' only require a single subtraction operation.

2.6.2 Filters

When filtering any signal, knowledge of all constituents of the signal must be known. Table 2.1 in the previous chapter presented a series of filters that have been used in published articles that were used to extract PPG signals. Whilst each of these filters would have successfully attenuated everything apart from the PPG (and in some cases, everything apart from just the PPG's fundamental frequency, from which the heart-rate can be directly extracted), other components of the PPG are also removed resulting in a waveform that is not an accurate representation of the PPG. By performing filtering that is this narrow, two problems will occur; firstly, as just mentioned, information about the PPG is lost; but more problematically, artefacts that pass through the same filter could result in a similarly looking signal. Because information is lost, the separation and/or detection of the PPG and artefacts is made considerably more challenging.

To illustrate this, a PPG recording and white noise were both filtered using the same filter as shown in Figure 2.7. Two points are of interest in the figure. Firstly, the filtered PPG, as predicted, contains very little of the original signal; the shape in particular is more akin to a sinusoid, the dicrotic notch has disappeared, and the leading and trailing edges of each pulse have become symmetrical. Secondly, by observation, the two original signals (the unfiltered PPG and the white noise) are nothing alike, however, the filtered signals are almost identical. This means that with this particular filter, noise that would otherwise be detected as such, would now look exactly like the filtered PPG: a potentially dangerous error.

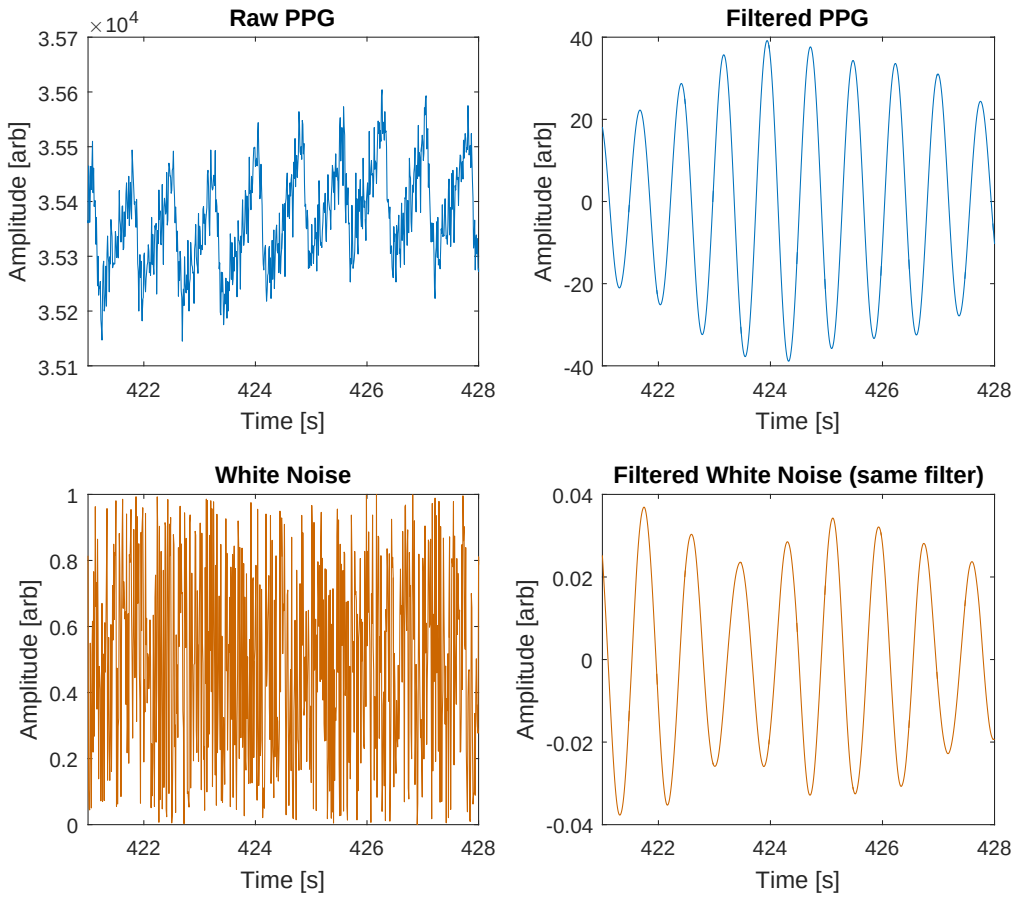

Figure 2.7: *Example of the problems associated with 'over-filtering' signals with narrow filters. Both the PPG and white noise are filtered with a band-pass filter, the same bandwidth (0.17 Hz) as Wu et al. (2012) [73] used.*

Figure 2.8 shows the same raw PPG and white noise as used in Figure 2.7, but with a wider, more appropriate, filter. Although some of the information in the PPG is still lost, the shape is now unmistakeably that of a PPG; the dicrotic notch is visible, and the rising and falling edges of each pulse are now distinctly different - the transition times are considerably different.

Both of the filters used in the previous illustrations are band-pass filters that in addition to high frequencies, also remove the low frequency components; this may not be desirable as the low frequency components may contain information pertinent to the PPG extraction. An example of why this is the case is presented in Figure 2.9.

If an artefact-corrupted PPG recording was blindly filtered with a band-

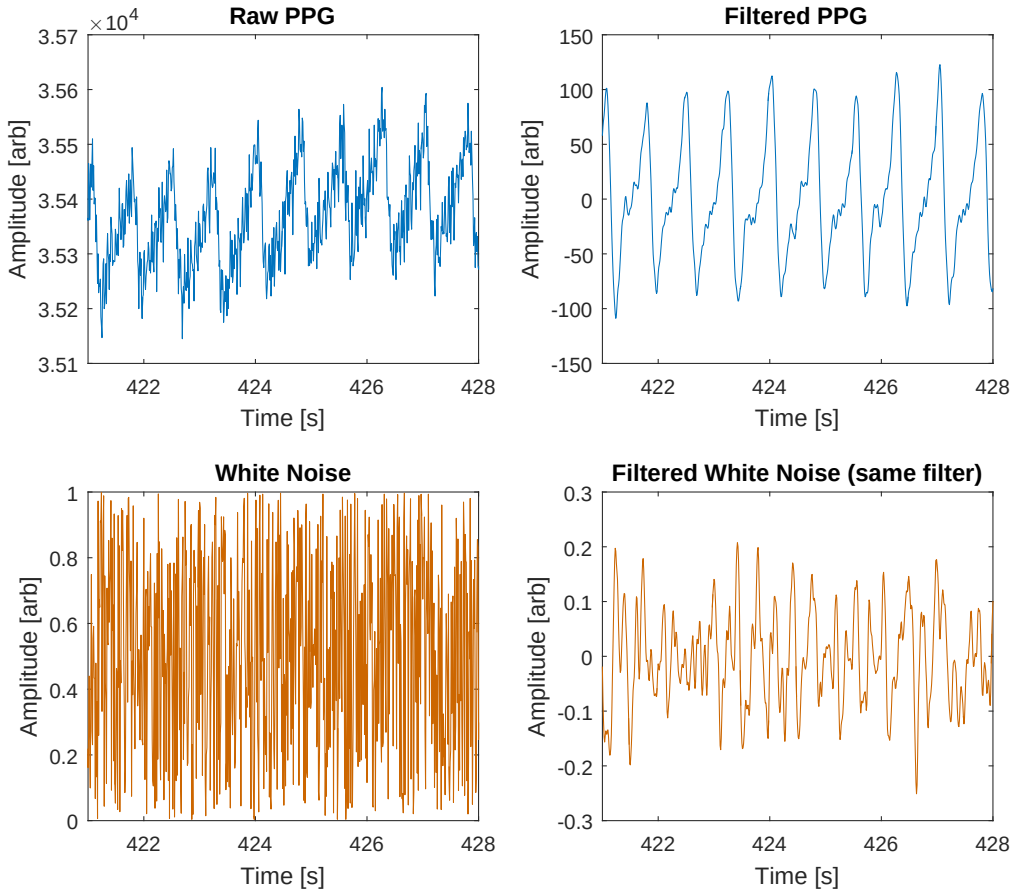

Figure 2.8: *The same data as in the previous figure, but now with a wider bandwidth filter (5.0 Hz). A clear difference is now visible.*

pass filter, so as just to keep the expected heart-rate frequency range (plus a few harmonics), then the resulting signal would contain the desired PPG, but also the artefact within the same frequency range. Analysing the frequency components of the signal would therefore not yield any meaningful information. In Figure 2.9 (right, centre), the corrupted PPG may still be used, but if the heart-rate extraction algorithm was not exceptionally good, it will fail. In order to determine where artefacts could exist, the out-of-band frequencies may provide the answer. In the lower right graph, the attenuated low frequency components of the original signal are shown to be of a high amplitude; this is evidence that suggests that the original signal was corrupt as amplitudes of this magnitude are not present in the 'normal' PPG.

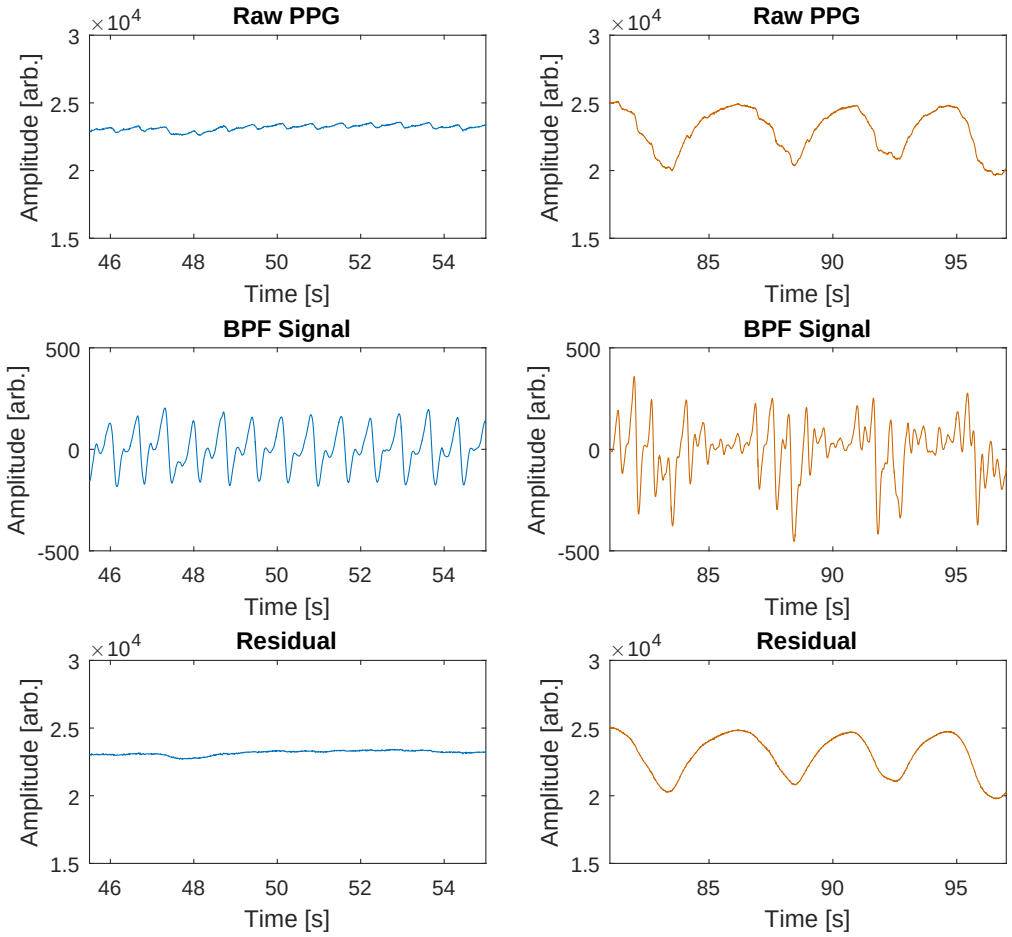

Figure 2.9: *Comparison of two band-pass filtered PPG signals, one without an artefact (left), one with (right). The central row shows the filtered signals, the bottom row shows the difference between the original (top) and the filtered: the residual.*

2.6.3 Frequency analysis

Cardiac-synchronous, pulsatile, signals such as the PPG contain a 'fundamental frequency' which is, in essence, equal to the average heat-rate over the measurement window. For heart-rate extraction algorithms, this is the only relevant information in the data. However, the pulsatile signal itself will contain many additional frequencies that give the signal its shape. For the PPG, the dicrotic notch presents itself primarily as a frequency that is twice the fundamental (a crude, but reasonable approximation of a PPG can be constructed from the sum of two sinusoids: $sin(x) + \frac{2}{3}sin(2x)$, Figure 2.10).

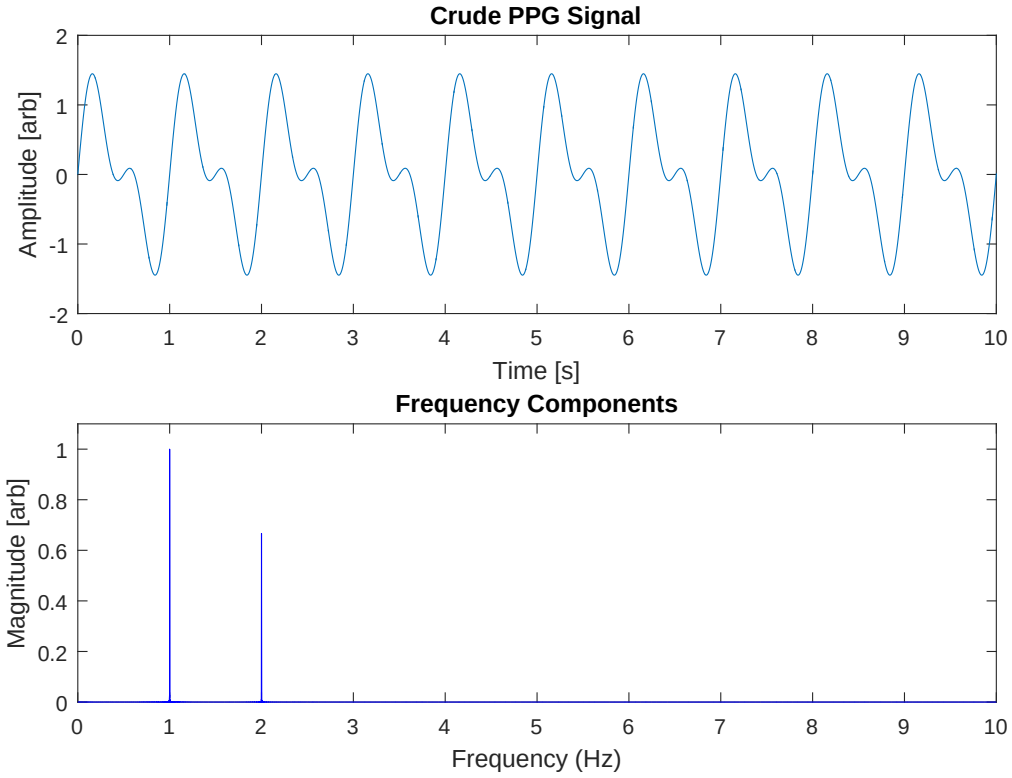

Figure 2.10: *A crudely constructed PPG and its frequency components.*

Whilst an approximation of the heart-rate can be obtained by using the dominant frequency from a frequency transform (a Fast Fourier Transform, FFT, for example), this must be computed from multiple pulses in the signal, effectively averaging the heart-rate over the window. As the window size decreases, the resolution of the FFT, likewise, decreases. There is therefore a balance to be made between responsiveness and accuracy (temporal resolution and frequency resolution).

Frequency analysis via transforms can be useful when analysing the signals themselves when in controlled conditions. This allows for generalised specifications to be constructed about the signals which may be useful when designing acquisition systems and processing algorithms.

As the PPG is a relatively 'stable' waveform (i.e. it's frequency, amplitude, and shape vary very little between different subjects, and the same subject under different conditions), accurate metrics can be obtained with a small set of data. Four datasets that were recorded at different times in this

project, for different reasons, and with different equipment, are presented in Figure 2.11, along with their FFTs.

Each of the four signals presented have been filtered with a first-order Butterworth band-pass filter. The lower cut-off was chosen to be 0.1 Hz so as to include the breathing/Mayer waves (whichever they may be) that are present in the signal, and the upper cut-off of 10.0 Hz was chosen so as to not visually attenuate any of the signal that is likely to be part of the PPG. Additionally, for illustrative purposes, the raw FFTs (in orange) have been smoothed to improve clarity. Vertical dotted lines are also included that represent the dominant heart-rate with two harmonics.

Whilst the 'crude' PPG reconstruction (Figure 2.10) provides a simple signal for testing basic processing techniques, these four FFTs illustrate the subtlety of the PPG signal and provide information about not only the 'bandwidth' of the PPG, but the relative amplitudes of the different major frequency components. Ignoring the lowest frequency component around 0.15 Hz (the breathing/Mayer waves), the PPG contains three distinctive frequency ranges; as stated before, the PPG fundamental provides the information about the instantaneous heart-rate and is dominant in the frequency plots; the other two ranges, at twice and three-times the fundamental, give the distinctive shape of the PPG. If these two components did not exist, only one frequency would be left, and the 'PPG' would look something like the over-filtered signal in Figure 2.7.

The frequency plot of the PPG from source 3 shows some distortion around the three components. This is due to a change in heart-rate and presents itself effectively as two overlapping FFTs, one for each heart-rate (this is visible in the unfiltered frequency plot, 2.11, Source 3, right, orange). This illustrates one danger of using an FFT to detect the heart-rate: with a window size that is too wide, the detected heart-rate may not be accurate, especially when the heart-rate is changing.

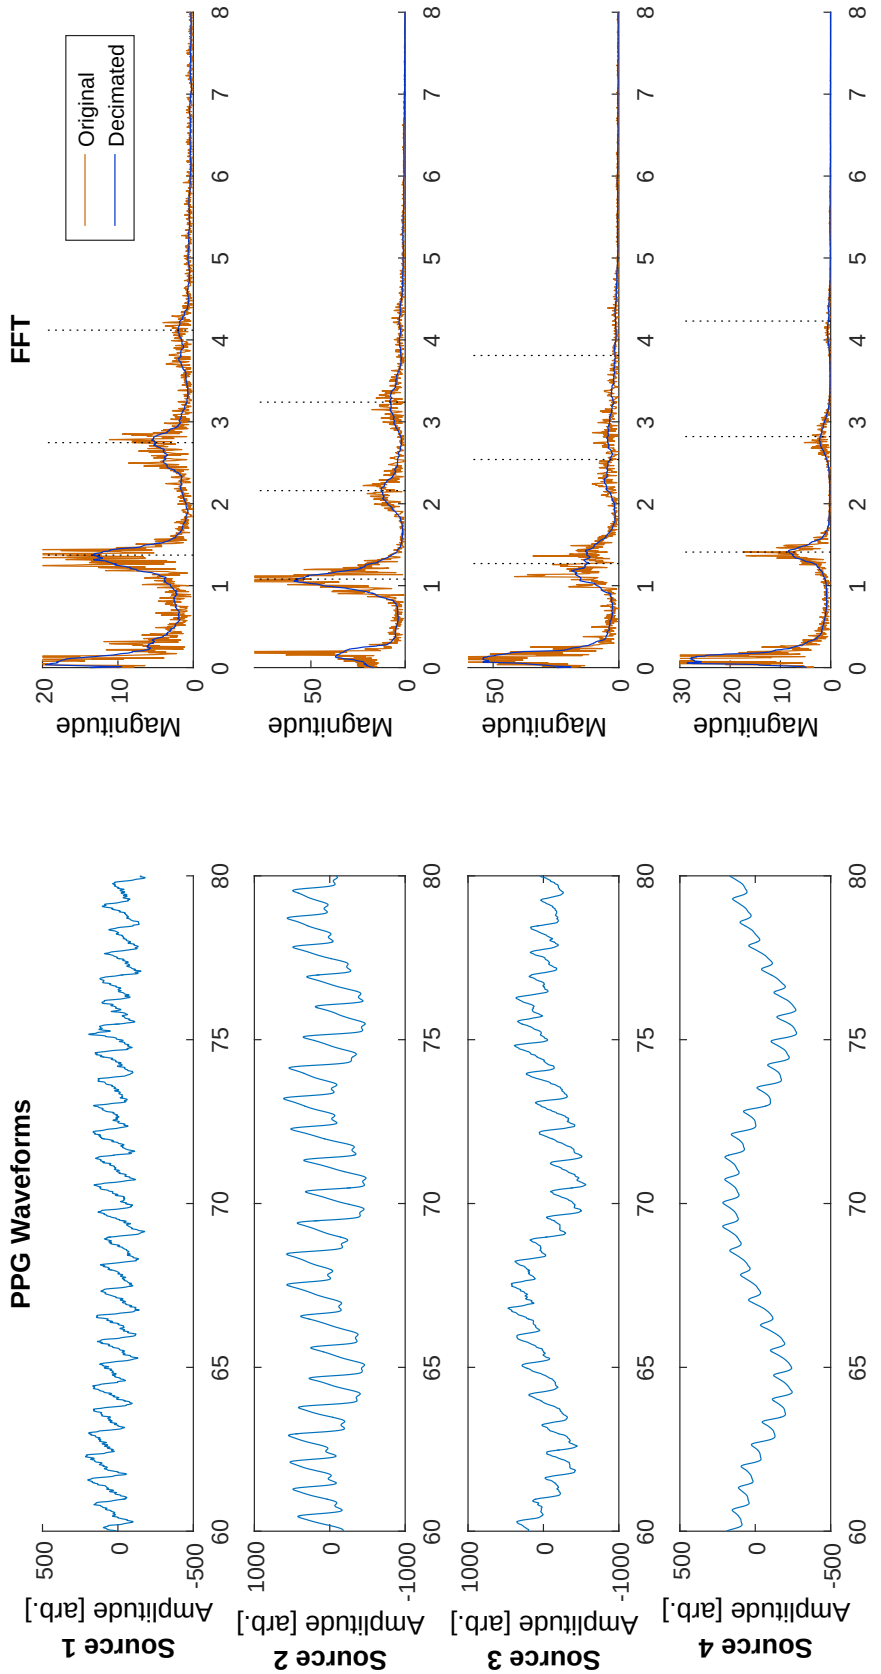

Figure 2.11: *Example of four 'good quality' PPG signals (see Chapter 4, for origins), band-pass filtered between 0.1 Hz and 10.0 Hz, with their FFTs.*

Given that the spacing between the three frequency components is equal, and typical heart-rate ranges are known, then the total expected band-width of the PPG can be easily calculated. This gives limits for the filter cut-off frequencies that will be used. Using heart-rate information from Fleming et al. (2011) ([81], page 15), the 1 to 99 percentile heart-rate range for newborns to 18-year-olds is from 40 bpm (1%-ile 18 year olds) to 180 bpm (99%-ile 1 month olds). To include the PPG's fundamental, plus the two harmonics, the absolute minimal bandwidth for a filter (to allow filtering of all age ranges) would be described by the lowest frequency, 40 bpm, and three times the highest, 180 bpm (giving 540 bpm). Filter cut-offs of 40/60 bpm = 0.66 Hz and 540/60 bpm = 9.00 Hz would therefore be required to ensure no loss of information for 98% of heart-rate ranges from newborns to 18 year-olds.

2.7 Discussion

Generally speaking, the research that has previously been looked at by other researchers has yielded promising results for their specific set-ups. However, for a general-purpose PPG acquisition system that is insensitive to motion artefacts, none of the systems would be appropriate. This is predominantly due to the strictly controlled environments where the tests have taken place.

Utilising some of the successful design choices from previous researchers (such as the choice of wavelength for the light source illuminating the skin), and avoiding the pitfalls that other researchers have fallen into (such as using filters that are far too narrow for extracting the PPG accurately or fairly), experiments can now be designed to test both aspects about the sources of artefacts, as well as systems being developed to more reliably extract PPGs from artefact-rich environments.

The following chapters present work to understand potential sources of artefacts (Chapter 3), ways to determine the effectiveness of PPG acquisition systems (Chapter 4), and a novel system design that allows for a more reliable extraction of PPG signals from a subject where motion is present (Chapter 5).

Chapter 3

Non-Contact Photoplethysmography

Imaging sensors (cameras) are useful tools for non-contact optical measurements, particularly as they are self-contained systems, easy to interface and readily available. They also allow for analysis over a large or multiple areas, and one or multiple subjects. The first practical work in this thesis involves a camera for those very reasons, used to test different aspects of remote (non-contact) PPG detection, and to look at reasons why motion artefacts are so prevalent in such measurements. Additionally, errors in the measurement technique itself (which will contribute to the artefacts) will be searched for. Ultimately, these exploratory experiments should provide illumination to possible methods of extracting PPG signals with lower sensitivity to motion.

3.1 Aims

As the work undertaken and presented in this chapter was initially exploratory in nature, the first aim was to gain a better understanding of possible sources of motion artefact. Afterwards, when hypotheses could be made regarding potential solutions to the motion artefact problems, the aim changed to constructing experiments to quantify artefact sources. Once such results were available, novel designs for PPG systems that have a greater tolerance for motion artefacts were generated.

3.2 Methods

To get a general understanding of the 'source' of the PPG: the skin (from the point-of-view of a sensor), a camera was set up to take several recordings of different parts of the body. The results were then analysed to get a better understanding of the skin's optical and topological structure; both of which, were hypothesised to be partly responsible for motion artefacts. This was the set-up for the bulk of the experimentation - where participants were to be recruited.

A few additional set-ups were made to illustrate other observed effects that arose from the primary experiment. These will be described and detailed when presented later.

3.2.1 Equipment & Protocols

A PCO PixelFly VGA camera (pictured in Figure 3.1) was chosen to perform the measurements, as it, at its core, had a 12-bit monochromatic CCD sensor. This allowed for much higher resolution data from each pixel (and with lower noise), such that the detection of the weak PPG was more likely.

With the camera operated with its maximum resolution of 640 by 480 pixels, and a frame-rate of 50 frames per second, the capturing system allowed for a maximum of 10 seconds of recording at a time - this was a limitation of the system (its memory, in particular) as the data was saved in an uncompressed format.

All recordings with the PixelFly camera were taken at 50 frames per second (fps), synchronous with the mains frequency (a synchronisation input was available on the capture-card, which was connected via an opto-isolator to a mains-referenced AC signal). This was to ensure minimal interference with the mains lighting (with power fluctuations at 100 Hz). This way each frame saw the same intensity from the artificial lighting.

Figure 3.1: *Image of the PixelFly camera used in the experiments. Image from [25].*

In addition to the room lighting, a single power-LED (3 W, $\approx$ 550 nm) was positioned on top of the camera in order to further illuminate the subject.

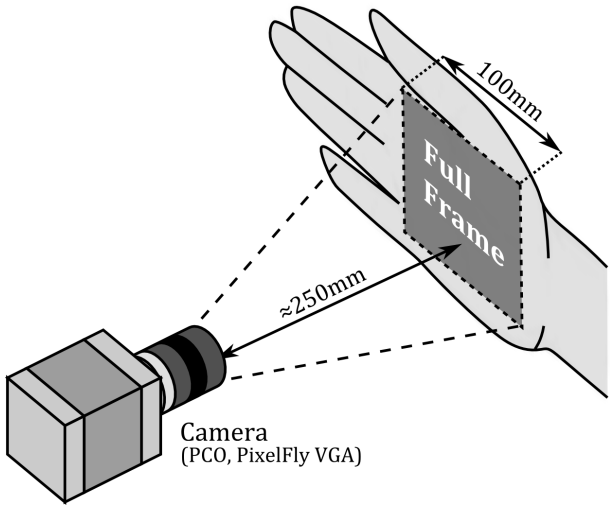

Figure 3.2: *Illustration of the arrangement of the camera and the subject. In this example, the palm is shown, but any part of the body described in the text is used. With an aspect ratio of* 4 : 3*, the size of the sensing region is 100 mm* $\times$ *75 mm.*

In total, six subjects were recruited for the trial (of mixed sex and ethnicity), and each subject has a ten second recording of their palm (whilst stationary) using the PixelFly camera, and images of their palm, the back of their hand, forearm, back of their arm, cheek (side of their face), and forehead, taken (set-up illustrated in 3.2). Each recording and image capture

was repeated three times. The University of Nottinghams Research Ethics Committee approved the trial (internal reference number is 2014140).

In addition, independent 10-second recordings were made of subjects holding an inanimate object at arm's reach. These recordings were made after an effect was observed during preliminary experiments. The reasons and set-up are described later.

3.3 Results

Note: Outcomes from the research in this section were published in the *Institute of Physics* journal 'Physiological Measurement', see [69].

Whilst many 'types' of movement exist between two objects (such as rotations (in multiple axes), 'zooming', etc), the intention was to use panning as it can be simulated after recordings have taken place.

Panning is the movement of the sensor's viewport (the image the camera 'sees') 'along' the subject, such that the distance between them remains constant. In cinematography, a pan is where the camera is effectively rotated in place, but from the relative point of view of the camera, the subject could also be moving in an arc around the camera. A reasonable approximation to this effect is to take a fixed subject and camera, and move the ROI within the frame. This is the technique that is used here.

To illustrate that a PPG was able to be obtained from the subjects, and in order to measure the amplitude (for later use), sample ROIs were used in the ten second recordings to extract the PPG. These are illustrated in Figure 3.3. At this stage, the smaller ROI is arbitrary in size and illustrates that by discarding areas that do not lie on the subject (such as the upper left corner of the image in the full-frame), an improved signal can be extracted. For this particular subject, a reasonable PPG amplitude is around 20 units (i.e., quantisation levels from the PixelFly camera).

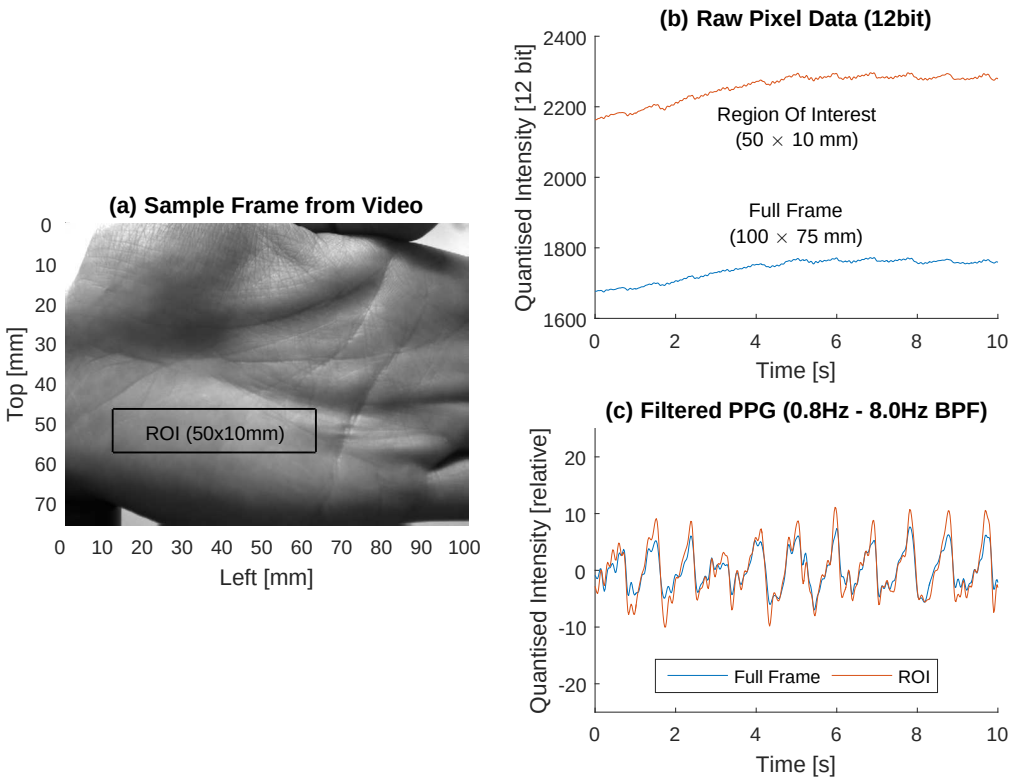

Figure 3.3: *Typical PPGs obtained from a small region of interest using the PixelFly camera, and the 'full frame'. The amplitude of the PPG (approximately 20 units) is used as a reference in the following figures. Bottom-right plot shows the two lines from the upper-right plot filtered with a first-order Butterworth BPF (0.8 Hz-8.0 Hz).*

To present the source of the artefact problems in as simple a way as possible, a single frame was taken from the recording (see Figure 3.4). In this frame a single row of pixels was extracted, and their intensities were plotted against their position. Additionally, the previously obtained 'typical' amplitude of the subject's PPG is also illustrated in the same Figure using a pair of parallel lines. Simply by looking at amplitude/intensity variations of single pixels of the image, it is clear that they dominate the PPG. Points (A) and (B) in this plot show the maximum continuous difference: a change of more than 600 units (30 times the amplitude of the PPG), in around 1 mm. It is clear that, if a single pixel ROI were to be used, then even the slightest movement would result in an intensity change an order of magnitude larger than the PPG.

By averaging the intensities from a vertical column of 200 pixels (centred

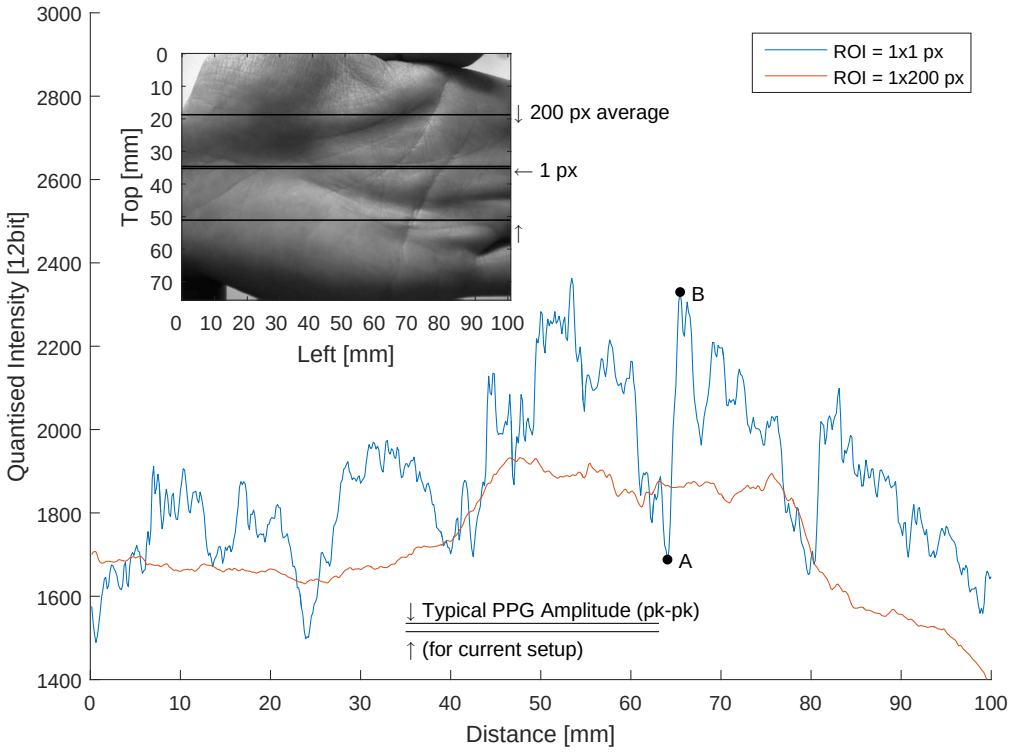

Figure 3.4: *Illustration of the intensity variations, as seen by the PixelFly camera, along both a single line and an average column of 200 pixels, on the palm. Even the averaged columns show considerably larger variations than the amplitude of a typical PPG. The 100 mm distance corresponds with the 640 pixels width of the image.*

at each pixel), an ROI with 200 times the area is formed. By plotting the intensities from these regions (also in Figure 3.4), a simple comparison of area is made. Although the higher (spatial) frequency components of the signal have been suppressed, the lower frequencies remain. The previously mentioned gradient (between the two marked points) no longer exists due to the larger ROI area.

Figure 3.5 shows the spatial frequency components of the two traces in the previous plot, alongside the measured amplitude of the PPG. Again, the lower frequency components dominate. By taking an average of multiple pixels (200 in this case), then the overall amplitude of the components is reduced (the higher frequencies more than the lower). This is an obvious result, but will be expanded upon later.

Of course, by moving the ROI linearly across the video (at a fixed ve-

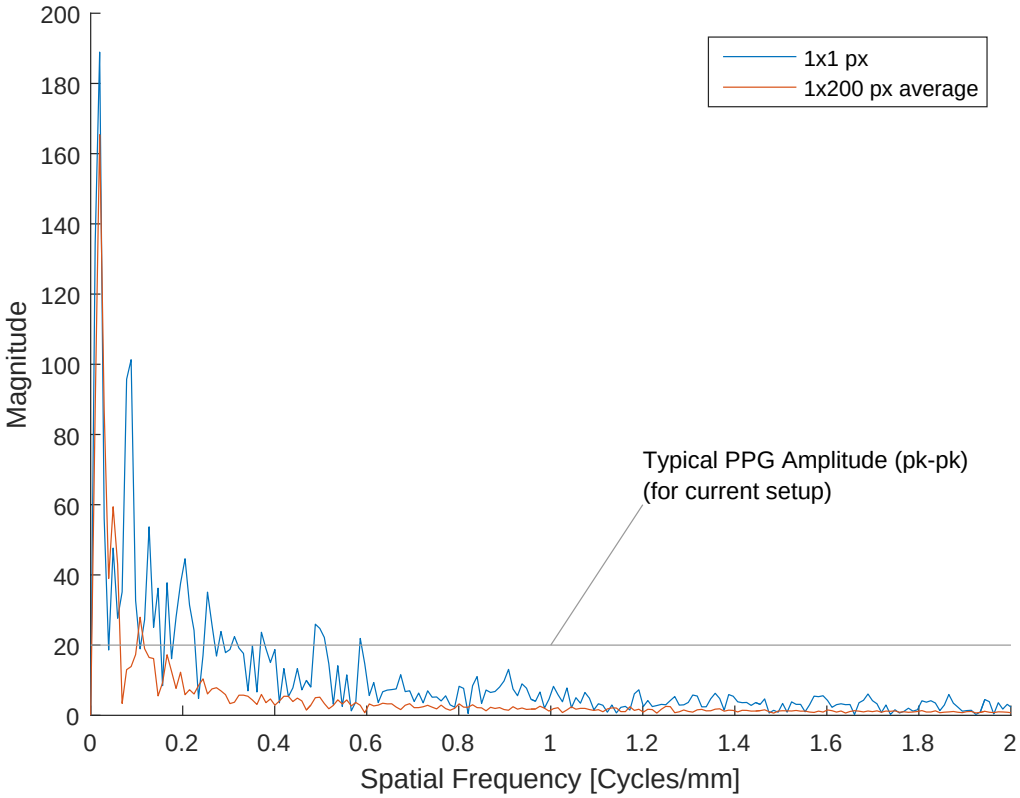

Figure 3.5: *Spatial frequency components of the waveforms from the previous figure (DC removed for clarity). This plot shows that the lower frequency components dominate the signal. The higher frequency signals are unlikely to cause a problem with the PPG.*

locity, or fixed number of pixels per frame), then the spatial frequency components can be interpreted as temporal frequency components, scaled proportionally by the speed. This is illustrated in Figure 3.6 where speeds of $1\,\mathrm{mm\,s^{-1}}$, $5\,\mathrm{mm\,s^{-1}}$ and $10\,\mathrm{mm\,s^{-1}}$ are shown. In each case, the spatial/temporal frequency components are just scaled versions of each other. At a speed of 0 mm/s (not shown), then the frequency components would collapse (be scaled) down to zero, and there would be no 'components' to speak of - no motion would mean no additional components (i.e., no motion artefacts).

Now it can be seen that the lower spatial frequency components, whilst not necessarily a problem at lower speeds, become more of a problem at higher speeds. This is due to them being 'converted' into higher *temporal* frequency components, which encroach on the typical heart-rate range.

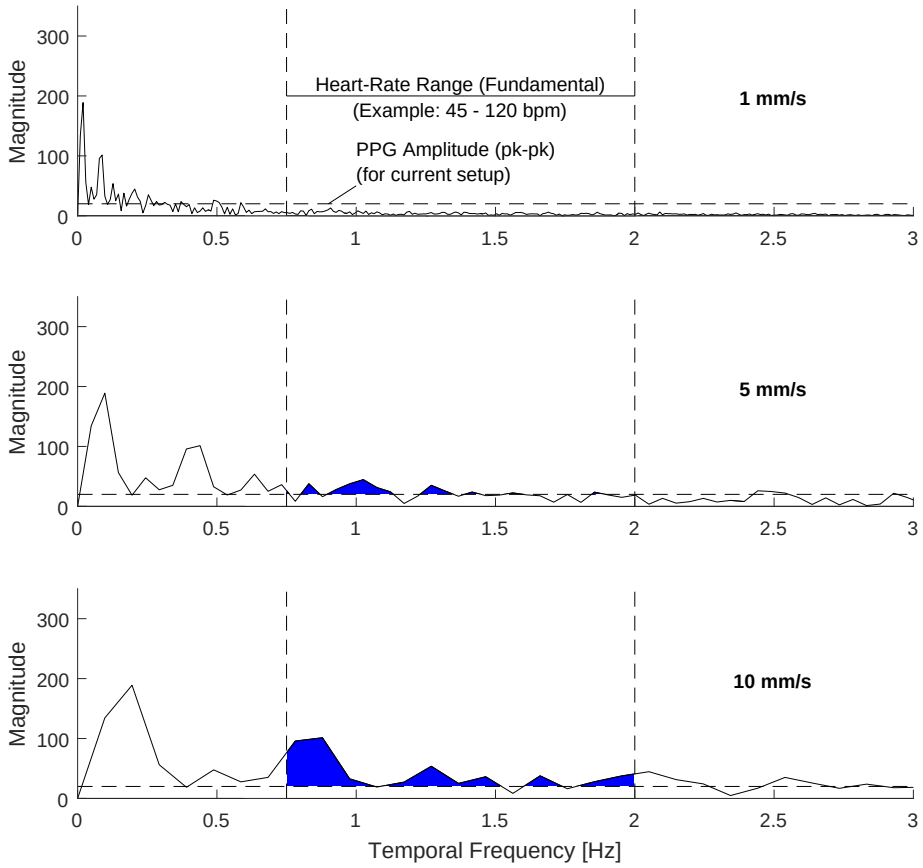

Figure 3.6: *Illustration of how the spatial frequency spectra (from Figure 3.5, 1 × 1 pixel ROI) of the subject's skin is transformed into a temporal frequency spectra when movement is considered. The lower frequency components (spatial), are more temporally dominant with greater speeds, and will affect the PPG more than higher spatial frequency components at lower speeds.*

Anything outside of this range can be easily filtered, but when the artefact overlaps this region, the process of heart-rate extraction becomes difficult - potentially impossible.

If an ROI was set to the largest possible area in a scene, whilst only covering regions of the subject where a pulsatile signal is present; and if the ROI was held stationary, then it can be assumed that the resulting signal will be close to optimum, regarding extracting the PPG. No motion means no artefact, and maximum area ultimately means a higher SNR. If the size of the ROI was now decreased to a smaller area, then as it is still being placed over the same subject (the same video), one could expect that the pulsatile signal would remain - with the pulses being in the

same position. This is true, but the SNR would now be reduced. Because the pulse frequency and positions should be identical in the two cases, a direct comparison between the two traces can be made. A simple Pearson correlation technique can be used to calculate a value that represents how 'similar' the two waveforms are. A value of '+1' means that the signals are identical, '0' means no correlation, and '-1' means that they are perfectly 'opposite'. By plotting the resulting correlation coefficient against the size of the ROI, a relationship can be made. The same technique can be used to measure the similarity between the signals when the ROI size remains constant, but with a motion present.

By repeating this method for multiple different sizes and traversal speeds, a 3D map can be constructed showing the correlation between each combination and the theoretical 'best' (of maximum size, minimum speed). This plot is presented in Figure 3.7. One important factor to note, that is not mentioned in the article published to Physiological Measurement, is that in order to get the ROI to move continuously for the ten seconds of the recording, without 'hitting an edge', when it nears an edge, the direction is immediately changed, such that the ROI 'bounces' between the limits; as the resolution of the camera was limited, this seemed like an acceptable means of extending the traversal path.

The shape of the plot in Figure 3.7 is as expected: increasing either (or both) of the ROI area or speed, decreases the correlation with the best possible PPG signal. However, the rate at which the two axes falloff can be calculated.

Figure 3.8 shows the same 3D graph as before, but viewed from one side (whereby the contours of the area axis are represented as separate lines). From this perspective, and with the help of lines fitted to some of the contours, it is evident that the correlation is negatively dependent on the square of the traversal velocity. The way the ROI area is related to

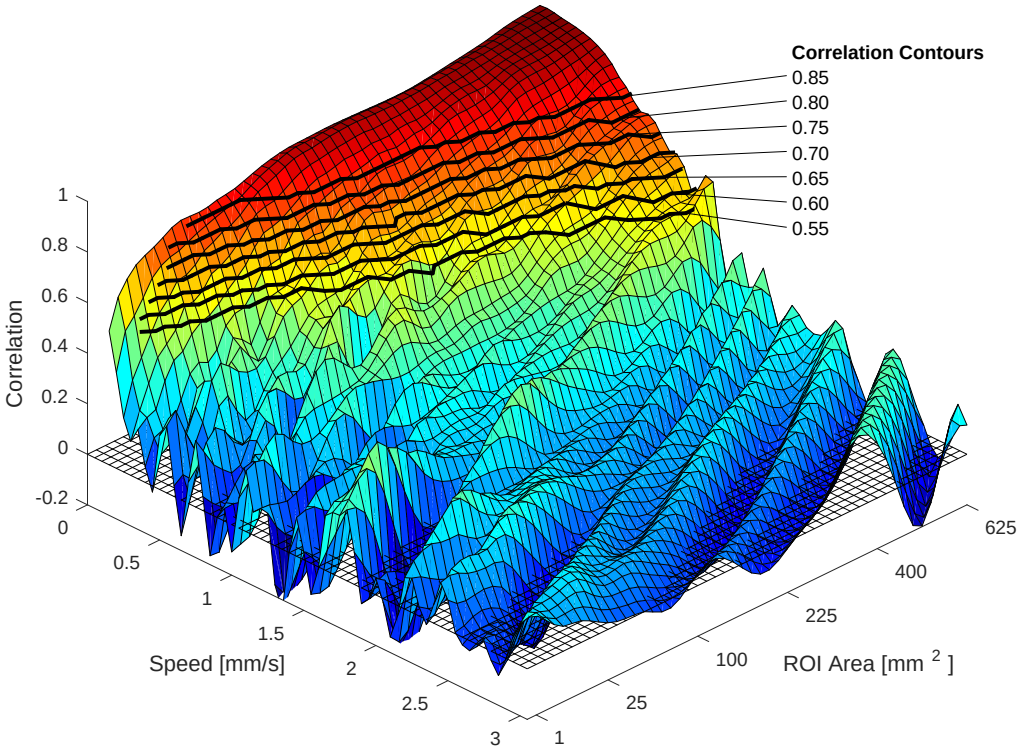

Figure 3.7: *Correlation of each ROI size and velocity variant with the 'best' case (of no velocity, and maximal area). The 'area' axis increases quadratically, to reflect how the area increases with the length of the sides of the ROI.*

the correlation is not so obvious, but can be obtained by looking at the coefficients of the fitted curves. This shows the correlation is negatively dependent on the reciprocal of the area.

By combining the two observations about the independent axes, a single equation can be formed that gives an approximation to the link between the ROI area, traversal speed, and resulting PPG correlation (see Equation 3.1).

$$Q = 1 - k \times \frac{V^2}{A} \tag{3.1}$$

... where V is the traversal velocity, A is the ROI area, k is a scaling constant that is dependent on the system (including parameters of the camera, light levels etc), and Q is the resulting 'Quality' factor of the PPG signal. In this equation, the quality factor is partially related to the

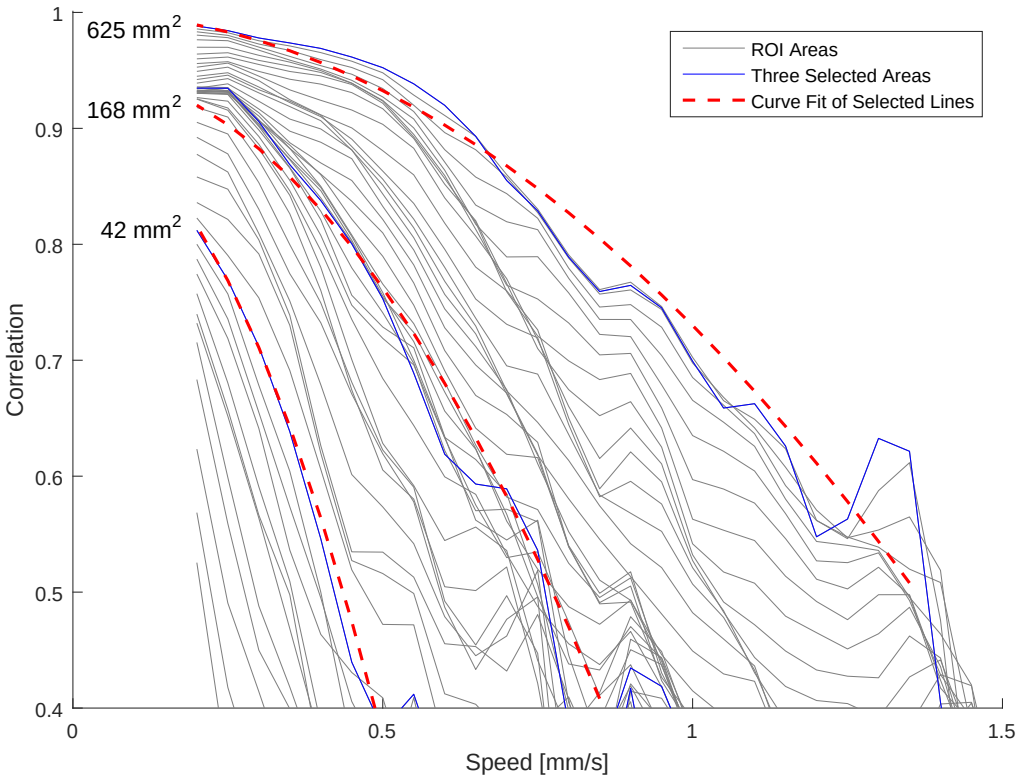

Figure 3.8: *Sideways 'slice' through the previous 3D plot, showing the falloff of correlation with speed (for each area). Equations for the highlighted lines (v is traversal velocity): (top)* $1.00 - 0.27v^2$*; (middle)* $0.95 - 0.75v^2$*; (bottom)* $0.90 - 2.10v^2$*.*

correlation coefficient, but is forced to a different domain of [0 1]. Here, a 'quality factor' of 1 is equivalent to a correlation coefficient of 1, where the PPG exactly matched the theoretical best-case scenario. A decreasing quality factor represents a PPG signal that increasingly deviates from this reference. A Q value of less than 0 is meaningless in this example, and should mean the same as a value of 0.

As this equation is dependent on the constant k, and that the constant is heavily dependent on the set-up of the system, this is not a universal equation that can be used between different scenarios. For example, in one set-up where quality factor is (correctly) calculated as being 0.5 does not mean that the same factor will be given for the same system but with the use of a different camera. Instead, relative comparisons should be made within the same system.

For example, if a Q of 0.8 is an acceptable value for a PPG when an ROI of size $42\,\mathrm{mm}^2$ is used, and a motion of $0.25\,\mathrm{mm\,s}^{-1}$ exists, then, in theory, the same Q can be obtained when the motion is doubled to $0.5\,\mathrm{mm\,s}^{-1}$, by increasing the area of the ROI by a factor of $2^2 = 4$ to $168\,\mathrm{mm}^2$. In the set-up as illustrated, this new ROI size is possible as it would still remain entirely within the bounds of the subject's skin. If the same doubling of speed were to happen again, the quadrupling of the area would mean that some of the ROI would cover an area that is outside of the subject's hand. This would result in a decrease in quality, and shows the limitation of the system.

3.3.1 Simulated Movement

A subject physically moving on their own is not easily reproducible, as it would be impossible to recreate the same movement 'pattern' each time. Needless to say the ever-changing physiology of the subject(s) will also change, making it difficult to make a fair comparison (one subject, with multiple recordings at different times, will rarely give the same result due to different heart-rates and movement).

If a ROI was chosen on a dataset where there was no movement, then the signal that would result would be a perfectly uncorrupted PPG. By artificially moving the position of the ROI each frame, then a 'virtual' motion would then exist. This motion (both speed and direction/path) can be accurately controlled and run on the same (or different) data sets each time.

Figure 3.9 shows an example of a recording from the PixelFly camera (10 seconds long, at 50 fps), and the resulting 'PPG' waveforms when the ROI is fixed at a single location (labelled 'fixed' in Figure 3.9), moving with a random direction each frame (labelled 'wandering' - movement of one pixel per frame in one of *eight directions*, spaced 45° apart), and wandering,

but with the same frame displayed (labelled 'path dc' - the same frame is displayed for each actual frame in the video). By taking the 'path dc', an *estimation* of the artefact can be obtained[1].

Given that each pixel in the images represents approximately 150 µm by 150 µm of skin (10 cm over the 640 pixels width of the frame, with equal aspect ratio pixels), then the average speed of traversal of the ROI, taking the longer diagonals into account, is given by:

$$framerate \times \frac{4 \times d_{vertical/horizontal} + 4 \times d_{diagonal}}{8} \tag{3.2}$$

$$50fps \times \frac{4 \times (150\,\mu m) + 4 \times (\sqrt{2} \times 150\,\mu m)}{8} \approx 9\,mm\,s^{-1} \tag{3.3}$$

In this particular case, even though the signal is clearly influenced by a motion artefact, the pulsatile component is not completely corrupted. It is still possible to extract the individual beats for some of the 10 second recording (particularly the first half). The second half appears to be corrupted slightly more, and it may be problematic to accurately extract the heart-rate.

As the wandering path progresses, it is clear that the *amplitude* (peak-to-peak) of the 'DC Path' increases, and is the cause of the corruption. By observing the path (lower right corner of Figure 3.9), it can be seen that there are two 'regions' that the ROI is focussed in. The ROI position starts at pixel coordinates (200, 140) and generally progresses upwards. When the ROI moves above Y = 130, the shadow that is caused by a crease on the subject's hand (in the lower right corner of the region) now lies on the edge of the ROI. When the ROI 'wanders' randomly, this high contrast region moves in and out of the area and causes the larger variations that are seen.

[1]Note that this is not the true artefact signal as the data from the subsequent frames is not used. This assumes that the only data that is contained temporally in the video is the PPG, which may not be true.

This subtle variation in the detected intensity of light is sufficient to corrupt the signal. If the ROI were chosen to be a little smaller (such that the high contrast region did not lie within the ROI), then the 'DC Path' amplitude variation would have been reduced a lot more than the PPG signal's SNR. However, reducing the size of the ROI comes with additional problems.

Figure 3.10 illustrates one problem that can occur when the ROI size is decreased. Compared to the 'fixed' ROI in the previous figure, the SNR of the 'fixed' signal in this figure is slightly degraded, but not sufficiently for the heart-rate to be irrecoverable. However, the effect that the reduced-size ROI has on the 'artefact-like' DC-Path is considerably different.

As the ROI is now focussed on a much smaller area, the same change in position can (and does) have a much more significant effect on the signal, as a result of the textured surface. The same high contrast region would incur a larger artefact due to the fact that a greater proportion of the region is affected. In this new example, with the smaller ROI, the 'DC path' gives a signal with greater amplitude fluctuations. When this is compared with the 'fixed' signal, it is clear why the resulting PPG signal from the wandering ROI is relatively more corrupted.

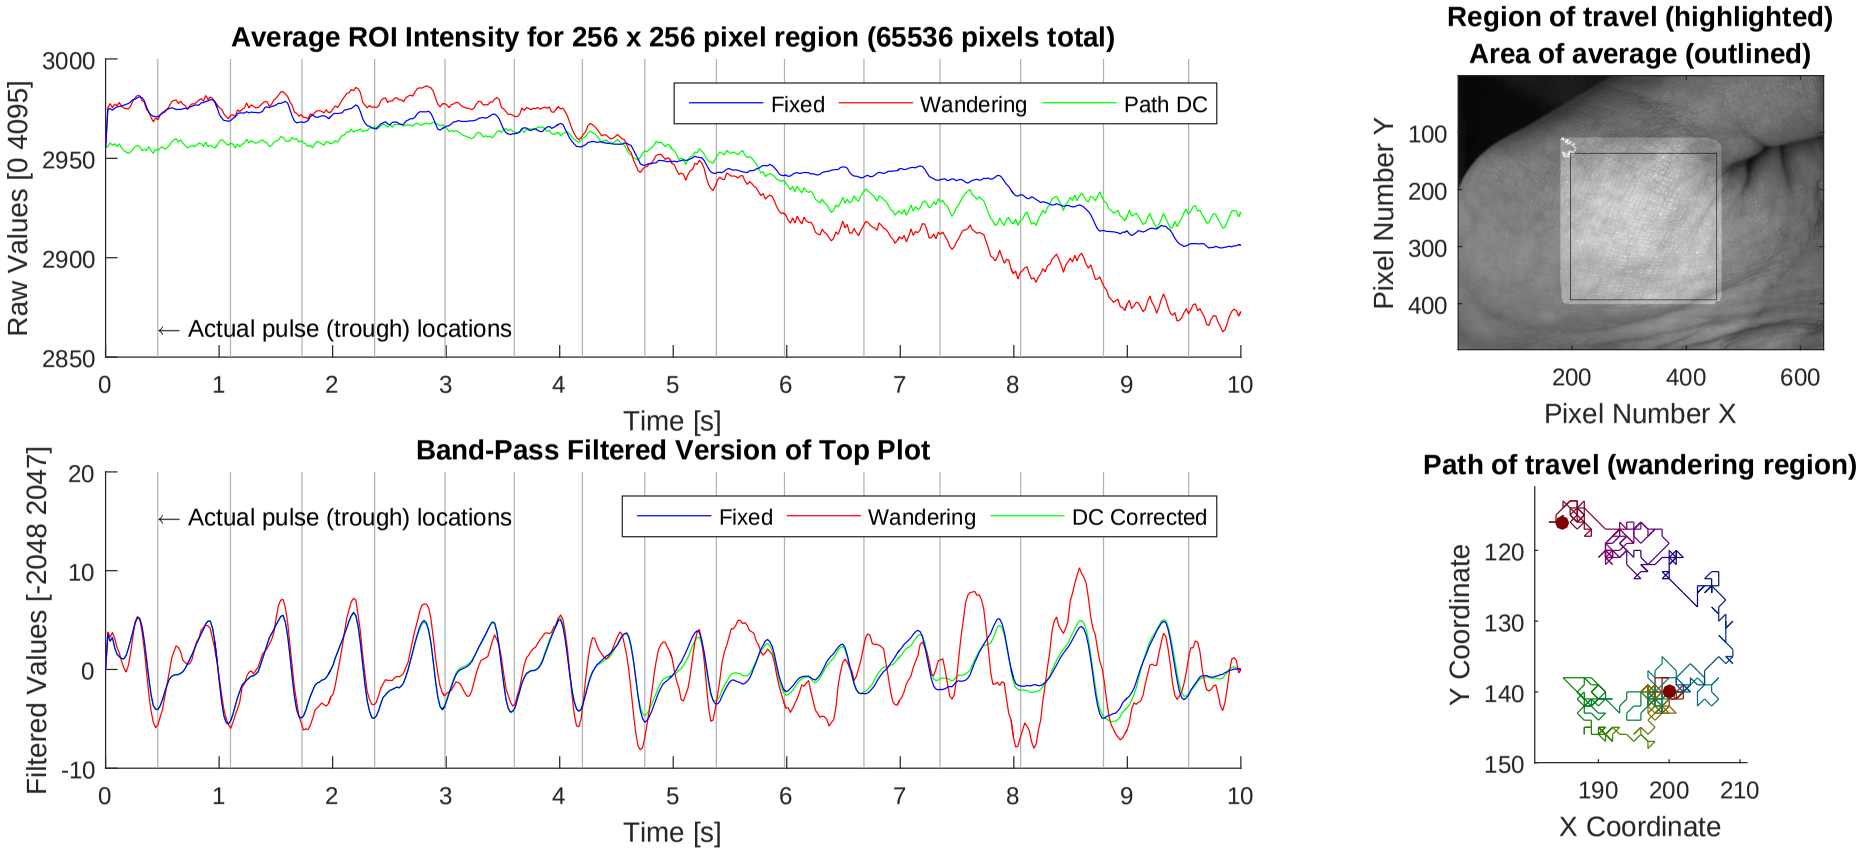

Figure 3.9: *The ROI was moved one pixel for each frame. Video of the hand, taken with the PixelFly camera. Top plot: the 'Fixed' line represents the PPG with a stationary ROI. The 'Wandering' line represents the PPG with a moving ROI. The 'Path DC' line is the moving ROI, but on just the first frame. Lower plot: filtered lines (second-order Butterworth BPF, 0.66 Hz to 9.00 Hz), apart from the 'DC Corrected' line which represents the wandering line with the DC Path subtracted.*

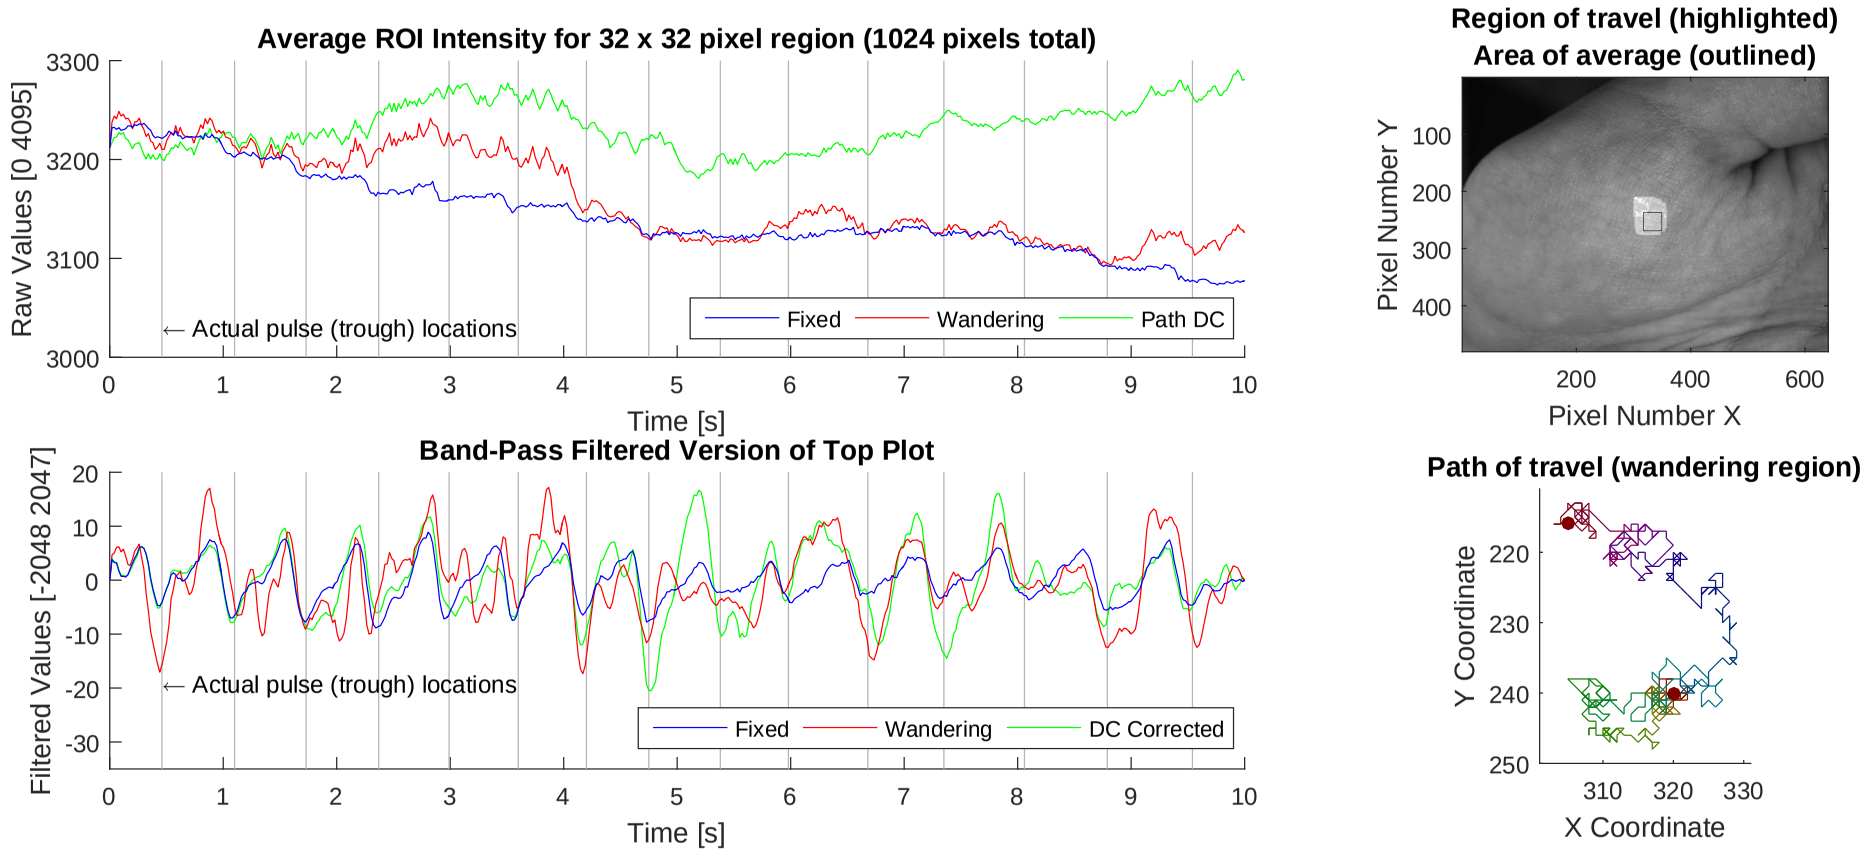

Figure 3.10: *The same wandering path as in the previous figure (on the same video, taken with the PixelFly camera), but with a smaller ROI of just 32 by 32 pixels.*

3.3.2 Cardioballistogram

Whenever the heart beats, the ejection of blood into the arteries has a mechanical effect on the body: it causes physical movement. The direct measurement of this mechanical movement is known as cardioballistography, and when the effect is observed via optical means, it can be referred to as optical-cardioballistography. This can either be a global effect on the body due to the acceleration caused by the blood's ejection, or caused by the expansion of vessels in the tissue resulting in a small (but detectable) movement; much in the same way as a hydraulic actuator operated (by 'filling' an expandable region).

For the second explanation (of the expanding vessels), this can be seen by simply observing movement of a long extremity - a leg, in this case - if it is under compression. By crossing one leg over the other, the large thigh muscles in the 'upper' leg are being compressed due to the leg's mass under gravity. The extension of this leg towards the toes acts as a large lever, amplifying any motion that should occur near the fulcrum. As the heart beats, the small expansion of the blood vessels in the thigh muscles and surrounding tissue causes a reciprocal effect on the toe/foot. This is observed as an oscillatory motion with a period equal to the heart-rate.

This effect was observed in three independent measurements of the PPG, one with a 'large scale' effect on the body (causing large movements, similar to the example explained above); one with *physical* movement on a much smaller scale, but due to the lighting set-up, causing a much larger optical effect; and finally, a direct observation of the palpations of an artery where it has passed close to the surface of the skin. These are each explored further below.

3.3.2.1 Large-Scale

In direct response to previously observed, pulsatile, **movement** of the hand during other tests, a simple, but interesting, experiment was used to test whether the heart-beat could be extracted directly from the mechanical movement of an arm. For this set-up, three measurements would all be taken synchronously with the same camera, one reference optical measurement (to obtain the background, ambient, light), one PPG measurement, and one ballistographic measurement (using a simple motion tracking procedure to measure mechanical deflections).

Figure 3.11 shows the set-up used. A (non-living) plastic baby, who cannot have a pulsatile optical signal associated with it, is held by a (living) human, who most certainly does[2]. The spatial tracking algorithm looked for the vertical position of the maximum contrast between the fiducial marker and the background in each frame. The other two measurements were of standard PPGs taken from the marked ROIs.

The expectation is that the (human) arm will move slightly with the heart-beat due to the variations in blood pressure within the arm's blood vessels. This movement should be detected both as a physical displacement, but also as an optical intensity variation (the latter caused by the former).

Figure 3.12 shows the signals obtained from the three measurement sites. As the *units* of the measurements are not consistent, the waveforms have been arbitrarily scaled for easier comparison. The upper plot shows the 'raw' values (with their means removed, for clarity), and the lower plot shows the same signals, but after a band-pass filter (0.6 Hz to 10 Hz, first-order Butterworth). The optical PPG and spatial signals both show highly correlated pulsatile components. The reference optical signal also has a somewhat periodic signal within it, but this is not correlated with the

[2]For this particular experiment, any inanimate object would have sufficed as a replacement. However, an artificial baby seemed appropriate given the aforementioned applications.

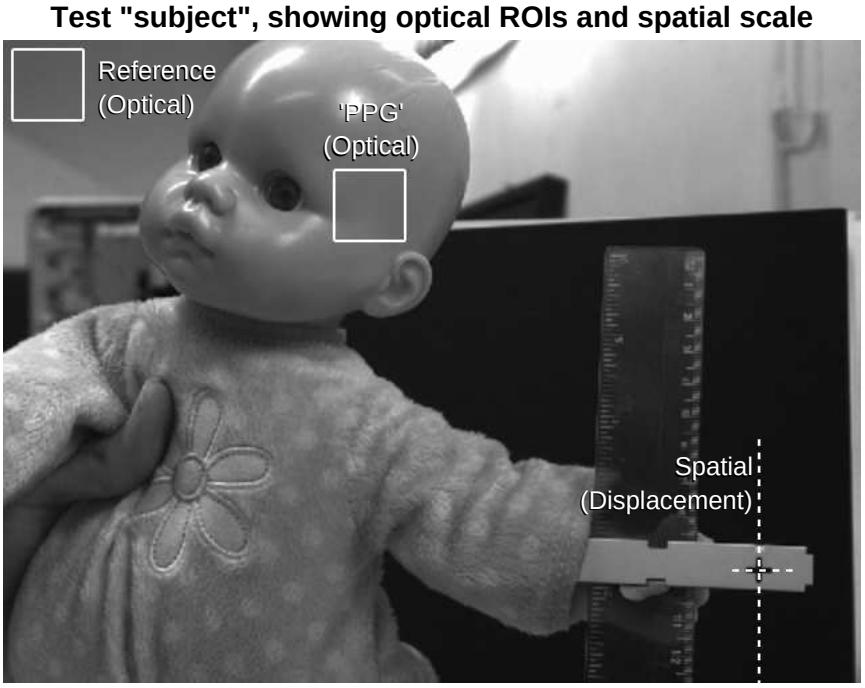

Figure 3.11: *Still image from a PixelFly CCD video showing the "test subject" (a plastic doll), being held at arm's length with a reference fiducial attached to its left arm (metal bar with a black cross). The dotted lines for the spatial measurement represent the detected position of the fiducial.*

other two waveforms (additionally, the scaling of this waveform is higher to make the three of similar amplitude - see the upper plot).

Whenever a non-contact PPG measurement is being taken, extra care must be taken to ensure that the signal that is being received is actually an optical pulsatile signal; this is due to the fact that whilst the signal *may* be cardiac-synchronous like the one shown in Figure 3.12 due to the physical effects of the heart-beat, it may also have origins from other mechanical sources: motion artefacts. When the motion is cardiac-synchronous, the mechanical effect on the optical signal is often of a much higher amplitude than the raw photo-plethysmographic signal, resulting in an often *correct* 'PPG' signal that is suitable for heart-rate monitoring, but a nuisance for when other statistics may be wanted (a cardioballistogram would not result in correct pulse-oximetry calculation, for example). Additionally, the delay or polarity difference between the optical and mechanical signals may be sufficient to mask the PPG entirely, resulting in degraded waveforms.

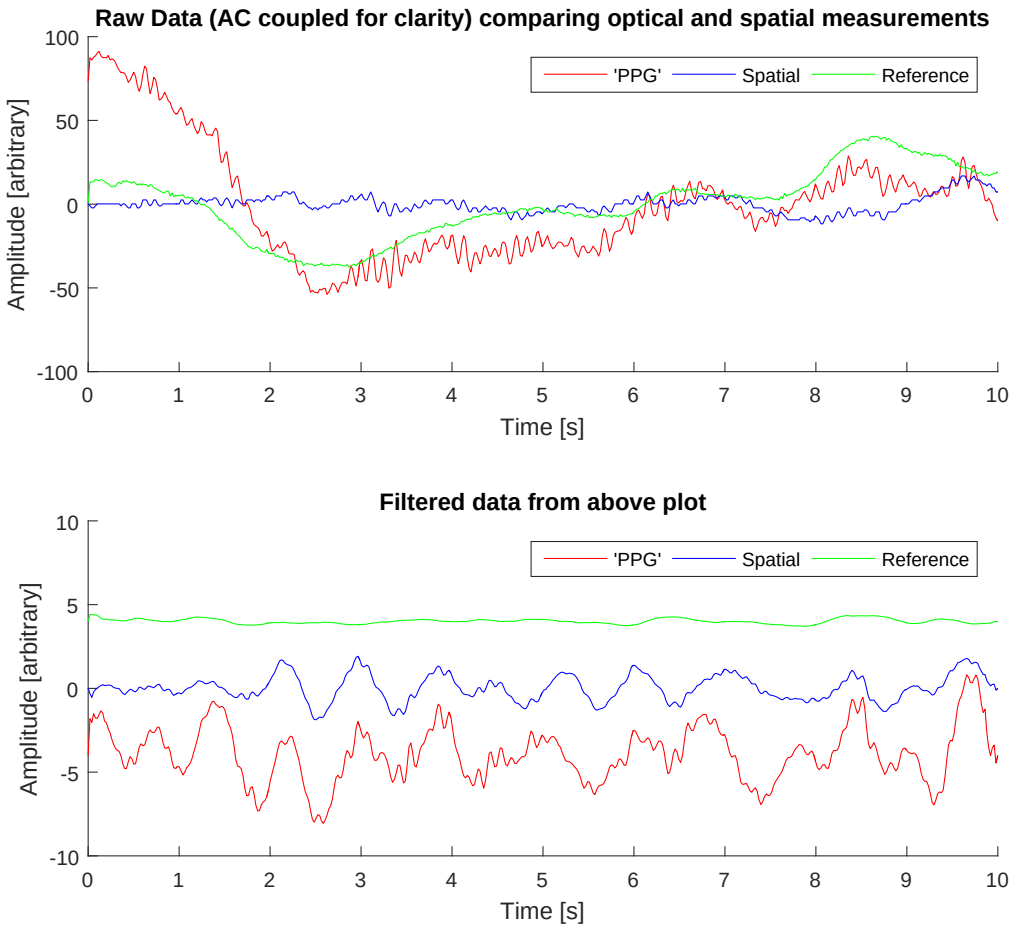

Figure 3.12: *Plots of the three detected values from Figure 3.11 of the optical reference, the optical 'PPG', and the spatial measurement. These plots have been scaled arbitrarily (vertically) for clarity.*

3.3.2.2 Small Scale

Whilst exploring how the angle of the skin (in relation to the camera) effects the extraction of a PPG, a curious effect was observed. Some regions of the subject's hand had an inverted version of the expected pulsatile signal. There have been mentions of the so-called 'inverted PPG' from other researchers [82], and each has proposed a reason for their existence. However, the work that the other researchers had published showed the inversions happening close to the light source and in the centre of the regions being investigated. The 'inverted PPGs' that have been observed in this research were located predominantly at the locations where the difference between the surface normal of the skin and the camera's direction

is maximal (i.e., at the edges of the hand).

To visualise the 'polarity' distribution of the pulsatile signals, each pixel was treated as a ROI and the waveform from that region was correlated with a sinusoid at the frequency of the known heart-rate (obtained via a separate in-contact sensor). The phase of the sinusoid was pre-aligned with the pulsatile signal obtained from the full-frame ROI; assuming the pulsatile signal from the full-frame image is a faithful representation of the true PPG, the correlation coefficient from each single-pixel ROI will show the polarity (if a correlation exists at all).

By plotting the correlation coefficient from each pixel as a colour (over one frame of the the video), a map is created of the distribution of the pulsatile signals' polarities. As a PPG signal cannot be reliably obtained from a single pixel of the PixelFly camera, *in theory*, the entire map should have fairly uniformly distributed colours representing no (or little) overall correlation. However, in reality, the map is far from this (see Figure 3.13).

It is clear that in this particular case, the generated map is not as it should be. There are very clear regions of highly correlated signals, either positively or negatively, that are **not** caused by the normalisation; if the correlation was uniformly low then the normalisation would leave the map speckled with highly contrasting pixels next to each other.

The three regions that are marked (with squares) were taken as secondary ROIs, and the waveforms that resulted from averaging the contained pixels are displayed in Figure 3.14. These regions were chosen to represent locations on the hand (in this particular case) where there were strong correlations with the reference sinusoid ('A' and 'C'), plus an additional ROI where there was no apparent correlation ('B').

The regions of highest opposing polarity appear on opposite 'sides' of the hand/thumb. More generally, locations on the hand that are 'facing' the upper left corner of the image show a negative correlation, whilst lo-

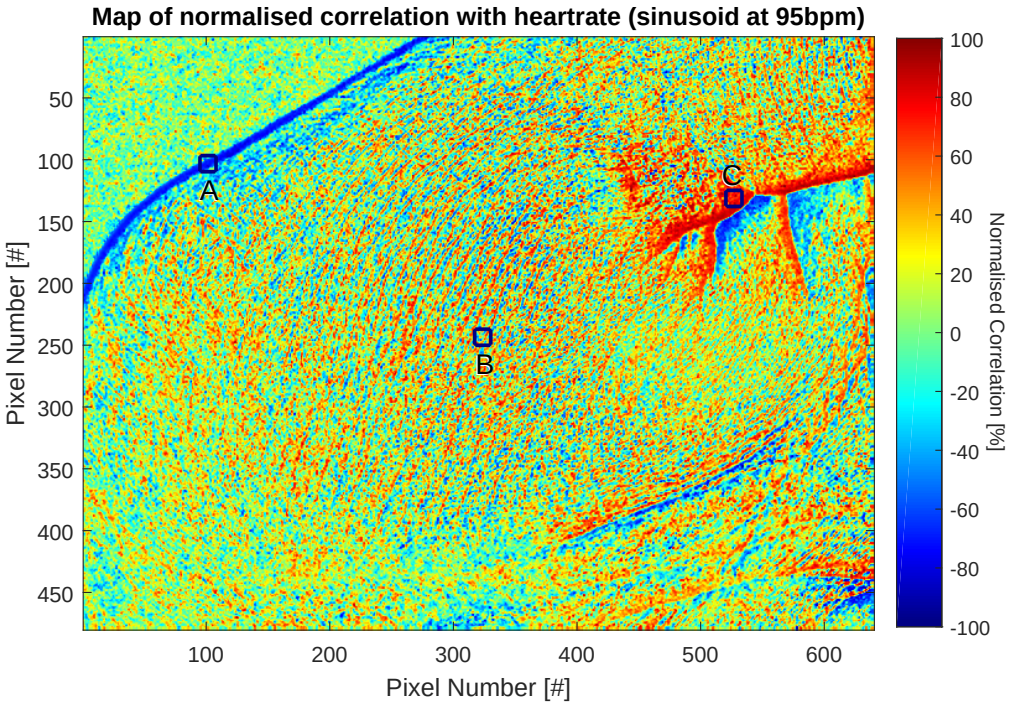

Figure 3.13: *Correlation map of each pixel ROI and a sinusoid with phase and frequency properties of the 'true PPG'. A normalised correlation of +100% represents a signal that is maximally similar to the reference ('C'). A value of -100% represents one that is maximally dissimilar (i.e: has opposite polarity, 'A'). A value of 0% represents no similarity ('B').*

cations with surface normals that face towards the lower right corner show a positive correlation. Small ridges on the skin, and a crease between the base of the thumb and the palm also show this distinct pattern.

This observation can be attributed to movement of the hand (a ballistographic effect); moreover, the movement is most likely a rotation. The axis of rotation is, from the point of view of the camera, from the lower left to the upper right corner of the hand. When a rotation occurs, one edge is turned towards the LED (that is placed on top of the camera), and the other, opposing, side is turned away. Each edge will experience an opposite change in intensity that is independent of the pulsatile absorption/scattering in the skin.

Overall, it is possible to see a slight bias in the map (in Figure 3.13) towards being 'positively correlated'. This bias represents the actual PPG, as caused by varying attenuation in the skin and tissue. If the whole image

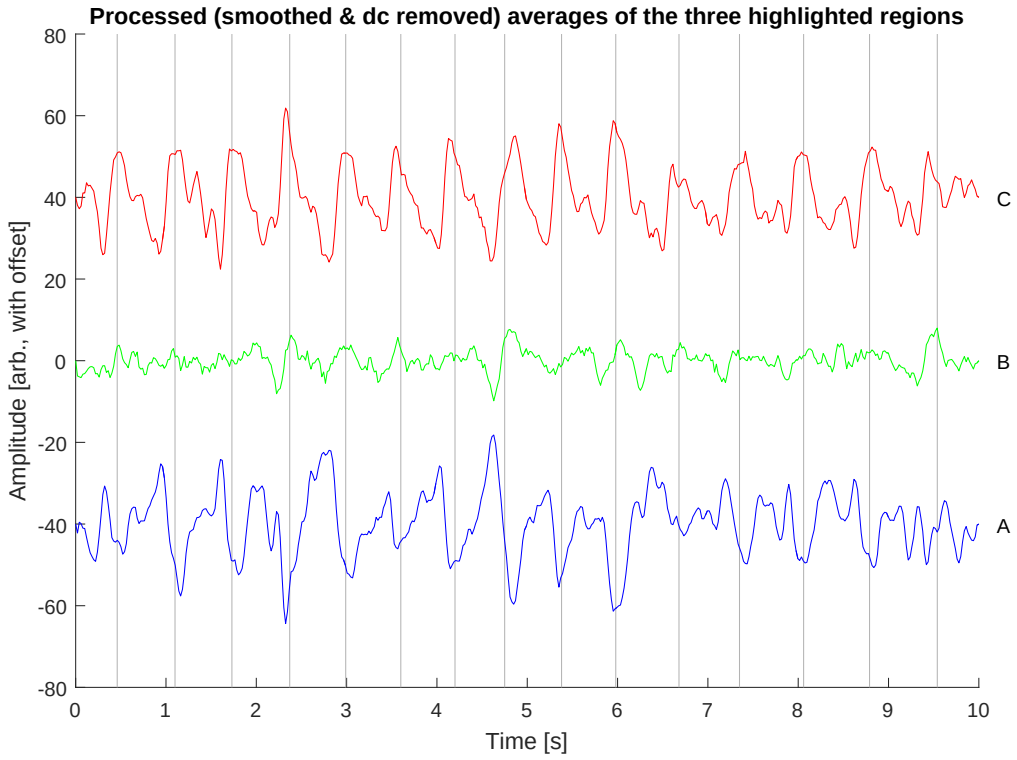

Figure 3.14: *Extracted signals from the three highlighted regions in Figure 3.13: the line labels correspond with the ROI labels. The red and blue lines are offset from zero to improve readability.*

were to be averaged as a single ROI, the resulting pulsatile signal would mostly be a result of the optical absorption; this is due to the 'positive' and 'negative' ballistographic signals effectively cancelling each other.

The important point to make from this result, is that by using a ROI that is too small, the localised ballistographic effects (not necessarily cardiac-synchronous) can easily dominate. A larger ROI would, as observed, take more area (that can contain opposing polarity signals) into account, resulting in a signal that takes *less* influence from motion. The issue that would still remain, however, is that even with a larger area being observed, a rotation of the hand would still have a global effect on the resulting signal if there was a light source not in-line with the camera's direction. In this case, a rotation would affect the pixels' intensity with a bias (more pixels would 'see' a change with one polarity than with the opposing polarity), leading to an artefact. For this effect to be removed, all non-aligned light

must be removed, and the total area of the part of the body being observed must be taken into account when detecting the PPG.

Note: This result is one of many that influenced the design of an in-contact sensor that reduced the effects of artefacts by taking, effectively, differential measurements of the system and combining them to remove their effect. See Chapter 5 for further details.

3.3.3 Polarisation

During additional preliminary experimentation with polarising filters, an interesting effect was observed. A single participant from the group has their Thenar region of their hand photographed with the same camera set-up, but with polarising filters applied to both the camera and light source. Two photographs were taken, one with the linear filters aligned, and one where they were orthogonal.

When using polarising filters in an optical system, there are two main effects: one positive, and one negative. The net gain (or loss) on the system will depend on how much each affects the system - whether the advantage outweighs the disadvantage. The primary *advantage* to using polarising filters is that particular components of the light that fall on the sensor can be partially attenuated in order to enhance the image. The primary disadvantage is that the filtering of the light will ultimately lead to less light falling upon the sensor. This attenuation can become significant, especially when the light levels that are involved are small to begin with (for example, when extracting PPGs).

In one particular set-up (shown in Figure 3.15), the polarising filter is aligned such that light from the hand that becomes polarised due to reflections off moisture/oils on the hand passes straight through the filter in front of the camera (due to their same polarisation), and be detected.

Both natural light, and the light from the LED (3 W, $\approx$ 550 nm, positioned on top of the camera, $\approx$ 25 cm away), are not polarised and will appear attenuated by this filter if *diffusely* reflected - half of the intensity will be absorbed by the filter. In this scenario, the specularly reflected light will dominate the scene (with diffusely reflected light being attenuated by half).

For PPG measurement, this is not ideal as the specular reflections will cause areas of high contrast due to the highly patterned surface. This is clearly seen in Figure 3.15 where ridges on the skin's surface appear particularly bright (upper-middle of this image). For small ROIs, this would be almost impossible to reliably extract a valid PPG from (during small movements), as the reflections would cause large intensity changes in the measurements. A similar problem exists for large ROIs as the intensity of a region will depend on the angle of the light source to the surface of the skin, creating 'hot-spots' with higher intensity - this is visible in the top-centre of the image. If movement were to occur with a ROI encompassing part of this region, then the detected value will change as it moved towards and away from it.

In Figure 3.15, an interesting effect can be seen in the lower left corner of the image. Here, the angle between the light source, skin, and the sensor is such that the polarisation angle is different. In the lower left corner, the polariser in front of the camera is attenuating the specularly reflected light. When the polariser is rotated, this region can be expanded (such that specularly reflected light is attenuated over the entire image - see Figure 3.16.

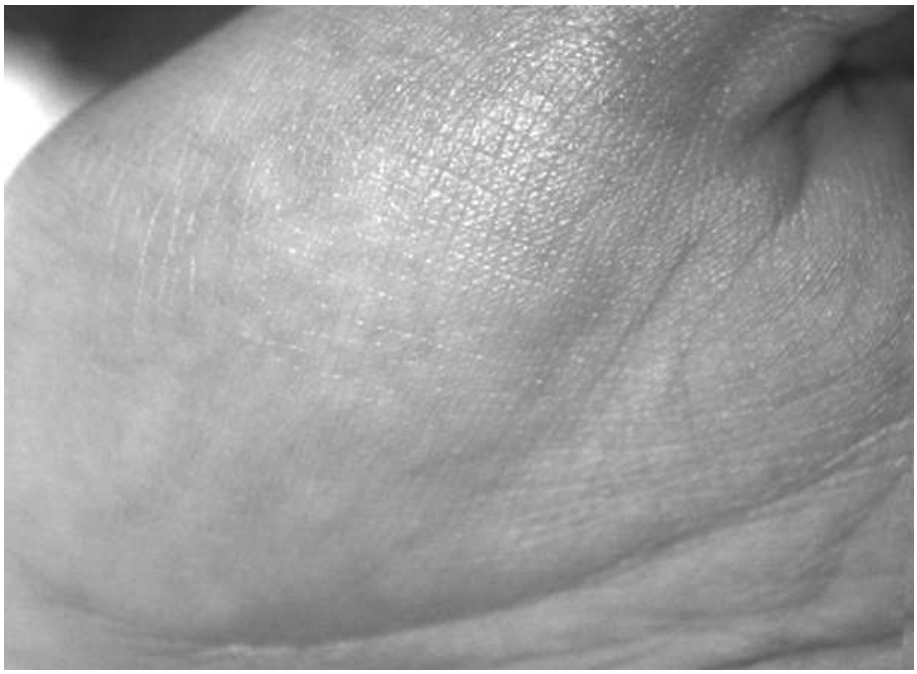

Figure 3.15: *Image of the thenar region of the hand, with a polarising filter on the PixelFly camera. The filter is rotated such that it has the same alignment as the specularly reflected light (they are 'in-line').*

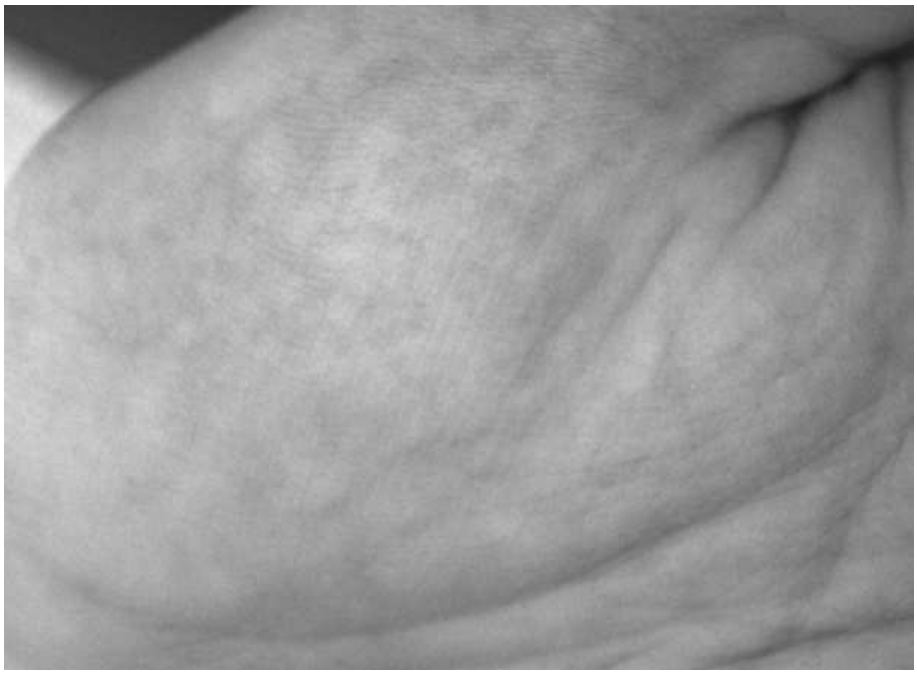

Figure 3.16: *The same region of the hand as the previous figure, but with the polarising filter rotated such that it has a 90° 'mis'-alignment with the specularly reflected light (they are 'crossed').*

In Figure 3.16, the opposite effect now exists. The specularly reflected light from the surface is now entirely eliminated by the filter, whilst the diffusely reflected light remains the same as before (attenuated by half). The diffuse reflections are caused primarily by scattering of light beneath the surface: the mechanism that allows photoplethysmography to work. By using a filter in this orientation (that is set-up dependent), motion artefacts can be attenuated. This is simply due to the fact that the light that 'emerges' from the skin is mostly due to sub-surface-scattering (where the PPG is 'formed'), and not due to direct reflections off the surface (where pulsatile signals generally do not exist).

Although less visible, at the location in the first figure where the large specular reflection existed, there still appears to be a region with increased illumination. This is a result of the diffuse reflection having *some* directionality to it. In other words, if the light source were to move, the location of the more brightly illuminated region would still change, despite there being no appreciable specular reflection. This would cause a motion artefact if one of the elements in the set-up (source/subject/sensor) were to move, although to less of an extent than before.

Although using polarising filters seems like a convincing benefit to the overall system, despite the attenuated light levels, there is one other factor that indirectly makes their use problematic. For the set-up as illustrated in the two figures, where the camera is directly facing the surface of the subject[3], and the light source is likewise aligned, an impossibility now exists whereby the attenuated specular reflection cannot exist over the whole image. In Figure 3.15, it was commented that the lower left part of the image had its specular reflections attenuated (just like in the following figure), due to the alignment of the filter. If the filter were to be rotated, then a different part of the subject's hand would have the same effect. The

[3]Such that the surface normal that is directly aligned with the camera's view direction is at the centre of the image

reason why the second figure (Figure 3.16) shows complete removal of the reflection is due to a different placement of the LED light source, such that the angle between the source, subject and sensor, was larger. This meant that the specular reflections would all have the same polarisation (as seen by the camera) and could be compensated for. By restricting the positions of the elements in the scene, if the specular reflections were not wanted, the usefulness of the non-contact imaging is deteriorated (simply due to there being some positions where the technique would not work - and would in fact make the resulting signal worse).

3.4 Discussion

When there is no mechanical coupling between the subject and light source/sensor, motion artefacts are much more likely to occur. This is simply due to the region of interest (ROI) being able to move over the surface of the skin, a material which is by no means uniform in colour or texture. Whilst some techniques are able to reduce the apparent 'variations' in reflected light intensity, such as using polarising filters, they require highly controlled set-ups that are not necessarily useful, and often, in-contact systems would be more appropriate.

Some measurement methods, such as using modulated light, apply to multiple applications in optical measurement. However, whilst a controlled (modulated) light source would greatly minimise the effects of ambient light in the vicinity of the measurement, it will have no effect when the subject moves due to the changes in specular reflection off the skin's surface.

One additional problem that exists with remote (non-contact) photoplethysmography is the fact that the heart-beat can manifest as an additional mechanical effect on the body (a cardioballistogram[4]), leading to

[4]Cardioballistography is, by itself, a separately studied field, and so its effects *can* be beneficial, instead of problematic.

uncertainty in whether the measured signal is a PPG or not. Whilst a ballistographic signal, measured optically, *could* reinforce the PPG signal, due to phase and polarity differences, this is not always reliable.

Ultimately, at this time, due to some of the limitations of non-contact photoplethysmography (as published: [69]), designing a system that has an increased tolerance to motion has been deemed too challenging for the current technology. In-contact approaches are the more feasible option for when reduction of motion artefacts is the goal.

It has been shown that by increasing the area of an ROI and averaging, different 'polarities' of the same artefact can be 'cancelled'. The PPG signals, of course, will have the same polarity regardless of the motion; therefore, taking a 'differential' measurement is a potential way of reducing a measurements' susceptibility to motion artefacts.

As this chapter looked primarily at the colour and topologically sourced artefacts on the skin (movement with spatial, not temporal frequencies), there was no need for a direct quality measurement of PPGs, as the PPG was not the output from the measurement. Following on from this chapter is work that was completed in order to quantify the 'quality' of a PPG signal. This would then be used in the final chapter of practical work in order to verify the effectiveness of the developed motion-artefact-tolerant system.

Chapter 4

Determining PPG Quality

When new devices are developed that aim to match or exceed the performance of other devices or equipment, there must be a direct comparison made between their outputs. In the case of heart-rate monitors, the comparison is usually performed on the calculated numerical heart-rate values; producing statistics representing the 'similarity' (correlation), or conversely, the error, between the values. Bland-Altman plots are typically used to illustrate the similarity by comparing the means and differences of the two data sets [83].

For heart-rate monitoring (and specifically, the heart-rate calculation algorithm), a "gold standard" device that has been previously approved is used. This most often takes the form of an (FDA/MHRA) approved ECG monitor. If the new device's output can show satisfactory agreement with the previously approved gold-standard, then there is adequate reason for the new device to be used in its place (although for the device to be accepted for use as a gold-standard, it will require more stringent tests, in order to prevent errors being propagated and/or accumulated).

When a device attempts to calculate the heart-rate from a dataset (either previously recorded, or real-time), it will not have the luxury of an additional 'gold-standard' input to compare its results with: if it did, what would be the point of the device? Instead, the device may need to calculate the 'quality' of a signal on its own in order to determine whether to reject it or not for heart-rate calculations, or for some form of 'quality index' that

may be presented to the user.

This chapter focusses on techniques and algorithms that have been developed and used to determine the quality of a PPG signal, without the aid of external stimuli.

4.1 Aims

Two interconnected aims exist in this section of work. One being to generate algorithms that will be able to quantify the quality or 'goodness' of a PPG signal, and the other to obtain a series of recordings of PPG signals in which to test the algorithms. To test the algorithms effectively, there must be a ground-truth signal available as a reference. Once these algorithms have been designed and tested, they can be used without a ground-truth signal - as described above - for the final practical chapter (see Chapter 5 where a robust PPG acquisition system has been developed.

4.2 Equipment

PPG measurement systems already exist (in both consumer and medical 'grades'), as do ECG measuring systems. However, few are combined into a single unit (these are mainly medical grade units that are found in hospitals mounted to trolleys), and even fewer are portable. Combining two smaller, portable, units together, retrieving data from them separately, and then combining the measured data afterwards is fraught with problems. Synchronisation of the two data-sets is potentially the largest issue. Getting better than one second accuracy between the units would be difficult, and there is often no reassurance that the 'sample-rates' from each piece of equipment would remain 'locked' throughout a large recording (oscillators in the equipment can be slightly off their nominal frequencies, causing drift over long periods of time, etc). Additionally, devices that are bought 'off-

the-shelf' will not usually output the raw data, instead, a processed version may be all that is available. This causes additional delays in the system, such that alignment is required. If the data of one or both of the data-sets is corrupt with artefacts, then alignment would become exceedingly difficult. It may be possible to use a (third) device to apply known 'pulses' to each device to use as a reference; but this would need to be performed, at the very least, twice in the recording (one at the start and one at the end), but this would be tedious and could induce additional (mainly user) errors.

The easiest solution to these problems, and one that gives the most configurable and reliable set-up, is to create a bespoke device that contains both PPG and ECG measurement systems that are inherently synchronised. This also allows for the raw, unprocessed, data to be saved from the simultaneously sampled data. As the 'clock' source is shared between the modules, then the samples that are extracted will track together.

This section presents the development of such a recorder, and the data that was obtained from it.

4.2.1 Requirements

The main requirement for the proposed design was that the two primary channels (PPG and ECG) were accurately synchronised. Ideally, if the two measurements were available as analogue signals, then each could be sampled with a pair of ADCs at the same sample-rate with the same sampling triggers. However, for the purposes of ease-of-construction, and to allow for different acquisition mechanisms, synchronisation within one sample (in time) was deemed acceptable. For a sample-rate of, say, 100 sps, the two recordings would be aligned to within 10 ms of each other. This way, any minute delays in the sampling components (sample-and-hold, conversion, etc) could be ignored.

Additional features of the recorder included the ability to wirelessly

stream the sampled data to a PC for real-time observation, and at the same time store the data to internal memory (an SD card) such that the recorder could be used 'out-of-range' of a computer. Naturally, for a portable device, it must be battery-powered, with a sufficiently large battery life to span several recordings.

The chosen sample rates for the two 'channels' was chosen to be 100 sps. This was set due to a limitation of the radio transceiver used (which had a limited bandwidth). This was, however, suitable for both PPG and ECG heart-rate measurement.

4.2.1.1 Safety

For PPG measurements, safety is not a direct issue as there is no electrical connection to the users, and the light levels involved are small (below the maximum permissible power for the given wavelength and continuous use [78]). For the electrocardiogram measurement, as a direct electrical connection to the user *does* exist, extra care must be taken when designing the circuitry. In order to comply with standards that govern (amongst other requirements) the electrical characteristics of a system, and how they interact with humans (EN-60601-1), the isolation of the system must be considered.

As the device is battery powered, and all regulators are linear (not switch-mode, using inductors, that *can* fail, so generating a higher voltage), the maximum voltage in the system is that of the charged, unloaded, battery voltage[1] (4.2 V in this case). However, when the system is charged, from either a USB connection to a PC, or via a 'power brick' (mobile-phone-like charger), then a risk is introduced if the recorder is to be used at the same time. Some devices use the same connector for both charg-

[1]Technically, during charging, the maximum voltage is 5.0 V. This is still well within the SELV (Safety Extra Low Voltage) classification in compliance with BS EN 60335 (due to the isolation).

ing and measurement in order to ensure that only one of the two options occurs at any one time. For this recorder, however, it was the intention to charge it via a micro-USB connector (so it can be easily charged with existing cables), which does not have enough connections to share with the sensors; therefore a power isolater was used. This takes the form of a small module that conforms to the EN-60601-1 standard, and provides 5 kV of isolation at DC. This was used between the USB socket, and the rest of the system. The ECG interface that was used was a dedicated chip that was designed for safe ECG measurement. And, of course, since the data communications occurs over a wireless connection, there is no electrical (safety) issue there. The chosen Bluetooth module is a Class II radio device with a 4 dB m maximum average output power (2.5 mW).

4.2.2 Design

It was decided (as an aside in this project) to develop a general-purpose development board that had all the necessary functionality to perform the tasks required for *this* project, but that could also be used by others on other projects with minimal development effort. The 'generic development board' would have a series of standard headers and connectors on it such that each project (including this one) could then simply develop a single PCB to interface with it. This would minimise development time and focus concentration and time on just the parts that are relevant for each system. Regarding the PPG/ECG recorder, a single PCB could then be developed that contained just the necessary electronics for the measurement.

To make the development board as useful as possible, it included the following features:

- Battery Powered (optional).
- Battery Charger.
- Mid-Range Micro-controller.
- ADCs.
- DACs.
- Wireless (Bluetooth).
- Real-Time-Clock (RTC).
- μSD Card.
- User Interface.
  - Buttons.
  - LEDs.
  - Speaker.
- External Interfaces.
  - UART.
  - I^2C.
  - SPI.
  - USB.

The resulting PCB (pictured in Figure 4.1 was made to be small and compact, whilst maintaining adequate accessibility to the external connections. A separate board was made for the aforementioned power isolator, for when the user does not need an isolated supply. The chosen micro-controller was a Microchip PIC24FJ128GC006, a 16-bit device, with a built-in 16-bit $\Delta\Sigma$ ADC, USB, and RTC (real-time clock).

For the specific purposes of this project, a small 'daughter-board' was made containing the system that is specific for the recorder. This is also shown in Figure 4.1 and contains, for the ECG measurements, a Texas Instruments ADS1298, an 8-channel, 24-bit ADC with an integrated 'ECG front end'. Although for the proposed experiments, only a 'three-connection' ECG is required, all eight channels that are available, including the 'right-leg-drive', are brought out for potential future use.

For the PPG side of the daughter-board, in order for multiple measurements to be made, two independent (and identical) channels were designed,

each containing a standard transimpedance amplifier, a *'bypassable'*, and programmable, band-pass filter, and a programmable gain amplifier (PGA). The outputs from which are directed to the ADCs. As for the LED drivers, the same general construction for the modules was used; but, since four DAC channels were available from the development board, four current sources were used. This allowed for multi-wavelength PPG measurements to be used on both of the sampling channels.

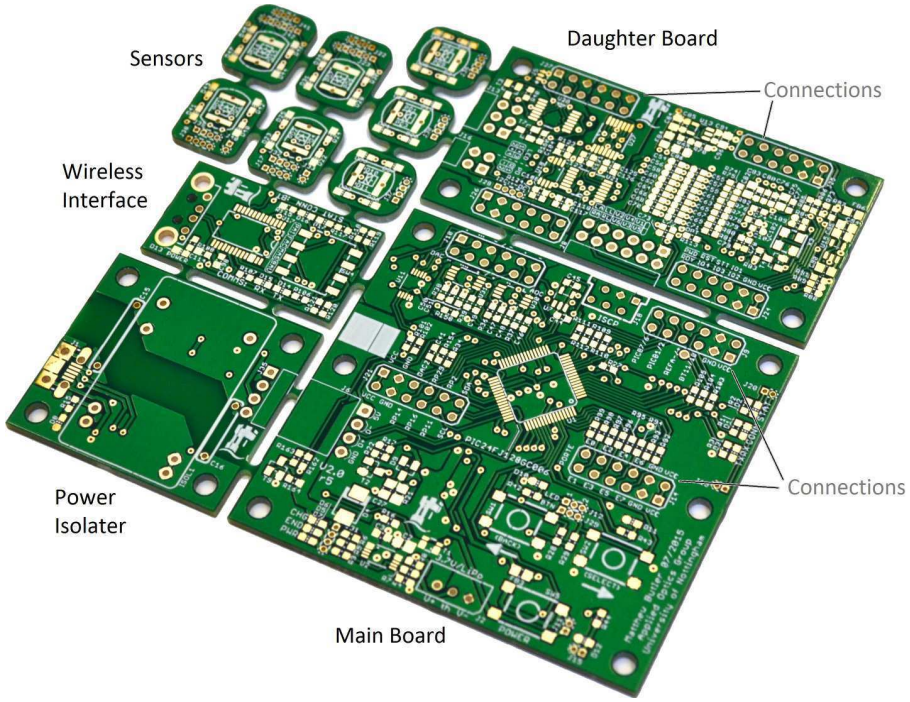

Figure 4.1: *Photograph of the unpopulated PCBs. The connectors in the daughter-board (indicated) line up with the main board (as do the mounting holes), allowing them to be easily connected together.*

Note: The full schematics and PCB layout for both the primary development-board and the secondary daughter-board are presented in Appendix II.

4.2.2.1 Completed Recorder

Figure 4.2 shows the completed recorders and associated cabling[2]. When developing the firmware for the device, the intention was to make the recorder have a simple 'terminal-style' interface, such that it could be operated using a command-line. However, as interest grew in the use of the recorder, it was decided (unnecessarily for this project) that a graphical user interface (GUI) should be developed.

The software GUI, illustrated in Figure 4.3 contains a real-time graph of the waveforms, and a pair of buttons (on the left) to start and stop the recordings. The second tab allows for direct access to the files that are stored on the μSD card within the box. The third tab is present for debugging and performing utility functions such as setting the real-time-clock and formatting the memory card, etc.

The files stored to the SD card are binary files and so require a program/function to decode them. As all processing was done in MATLAB, a short script was written to load the file 'into' a structure, ready for processing.

> At the time of the writing of this thesis, *five* individuals/groups have used either the development board or the recorder box(es) for their own projects, demonstrating the attractiveness of the customisable design.

[2]Three were made in total, each with a slightly different configuration, for other researchers wanting a recorder for their work.

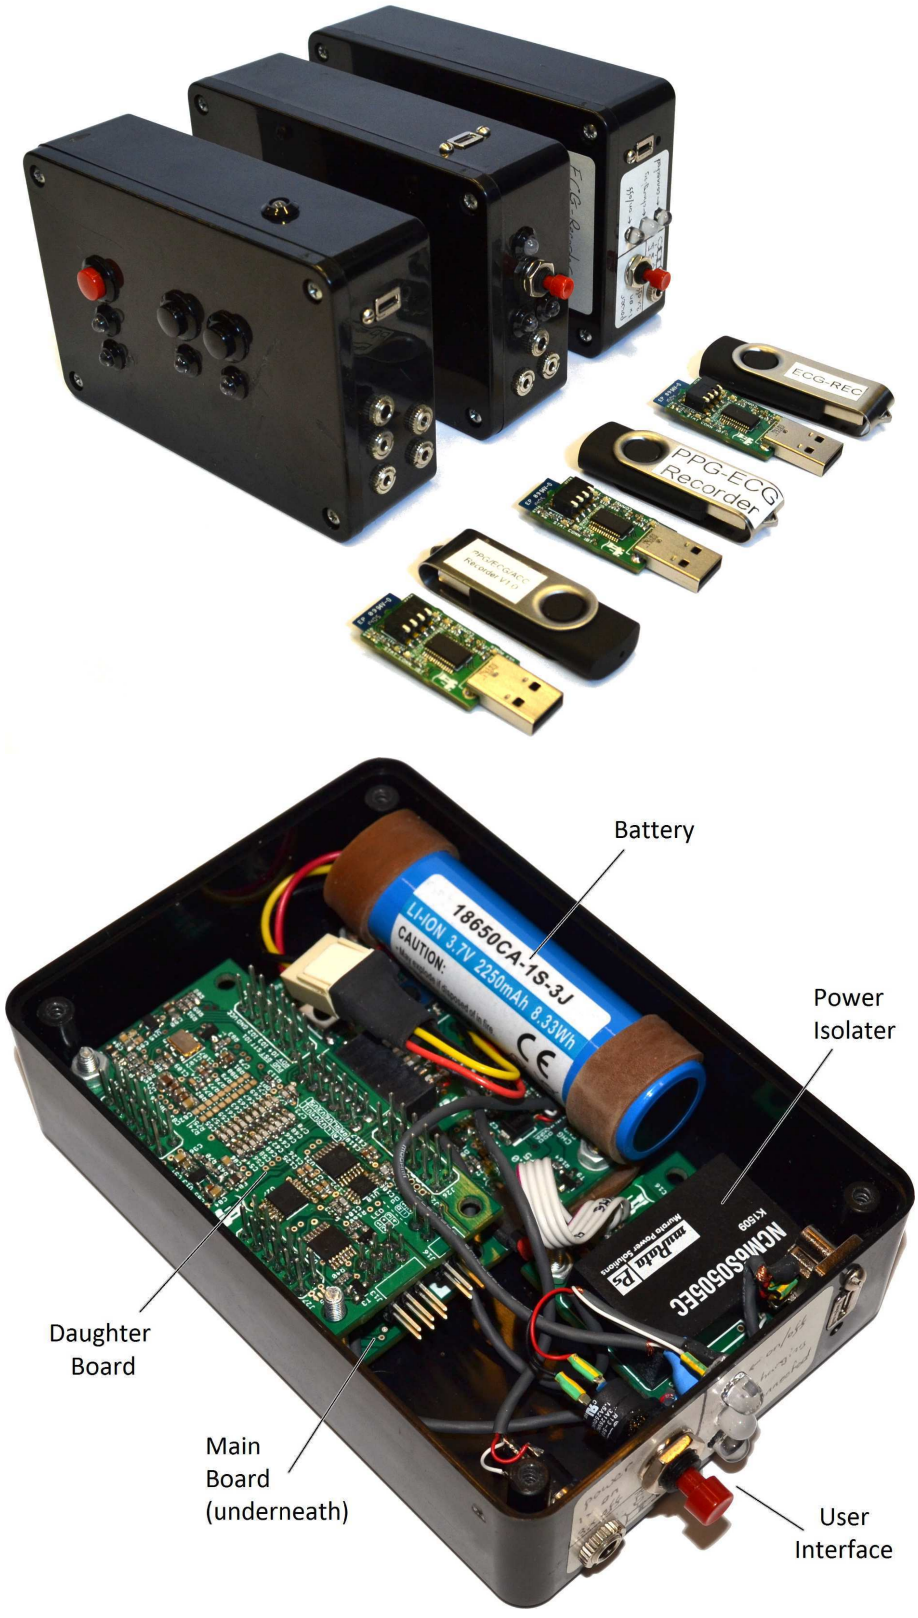

Figure 4.2: *Photograph of the finished recorders, wireless interfaces, and memory sticks containing the real-time viewer/recorder software and MATLAB scripts (top). Photograph of the inside of one recorder box's enclosure (bottom).*

Figure 4.3: *Screen shots of the software used to interface with the recorder box. The upper window shows the interface used to choose the recorder to connect to. The lower window shows the waveforms during a recording.*

4.2.3 Testing

Before any trials were run on participants, a very crude (but effective) physical simulation of a finger was constructed in order to test the effectiveness (and reliability) of the recorder. In order to simulate the absorption of the light within the finger, a monochromatic Liquid Crystal Display (LCD) was used as the absorbing medium. By quickly changing the characters that are displayed on the screen between a space (*no* pixels 'on'), to a black square (*all* pixels 'on') with a varying duty-cycle, the 'intensity' of a region on the

screen could be changed. By making the modulation frequency faster than the response time of the LCD, the transition between off and on could be made to be smooth (there would be no flickering).

Figure 4.4 illustrates how the LCD intensity technique works. The darker pixels absorb light that would otherwise be reflected off the rear surface of the LCD, behaving *similarly* to how variations in blood volume in tissue absorb the light. It should be stressed that this was **not** intended to be an accurate phantom of a real finger, more of a piece of test equipment that could be used to test the recorder in a *repeatable* fashion; the absorption caused by the LCD is fairly constant over different temperature ranges (unlike with a human finger, whereby the body shunts blood away from the extremities when cold), and can be set to have a constant frequency.

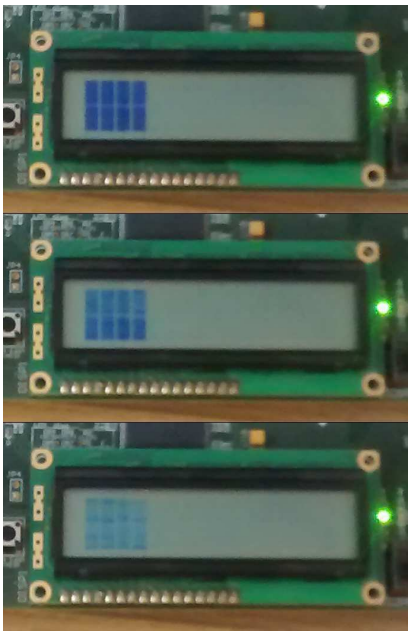

Figure 4.4: *Photographs of three different 'intensities' of the LCD pixels. The upper image shows the maximum intensity, the lower shows the minimum intensity (which is not completely 'off' due to the persistence of the display when modulating the intensity with a varying waveform).*

The 'pulse pattern' that dictated the duty-cycle of the modulation is illustrated in Figure 4.5, which includes a dicrotic notch. In order to compensate for the relatively slow response time of the LCD, the dicrotic notch

in the input duty cycle has been exaggerated.

Figure 4.6 shows the resulting waveform as acquired using the recorder box. The resulting PPG is 'inverted' from what is expected due to the amplification stage in the hardware being 'inverting' (i.e: with a negative gain). The value from the box can therefore be thought of as an 'attenuation' of the light, rather than the light intensity itself. To get the actual measured intensity, the value can be subtracted from the maximum measurable value ($2^{16} - 1 = 65535$).

The amplitude of the signal is relatively high compared to a PPG from a human, simply due to the large absorption variation from the LCD. This can be changed by moving the sensor head away from the LCD, or by placing a diffuse/translucent material between the sensor and LCD. The duty-cycle that the LCD is being run with can only realistically be set to one of sixteen different values[3], and so cannot control the absorption to a higher resolution.

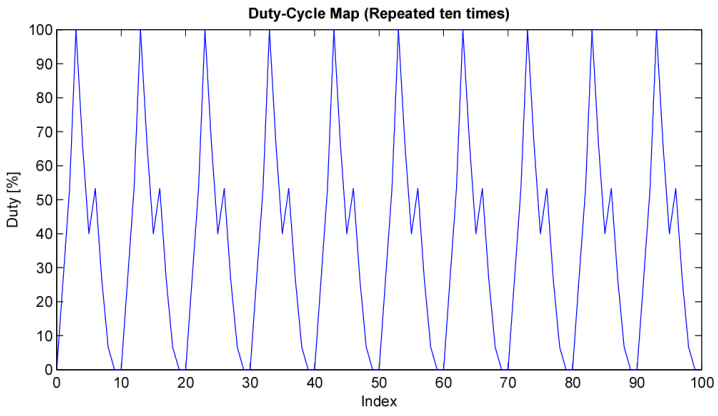

Figure 4.5: *A ten second sequence of 'PPG pulses' with an effective heart-rate of 60 bpm. The vertical axis represents the duty cycle of the modulated intensity.*

One final test for the recorder is to see how the recorded signals are affected by different LED brightnesses. If the electronics within the recorder incorrectly add an offset to the data (caused by the inverting amplifier's

[3]In order for the duty cycle to have a higher resolution, the modulating frequency must increase. However, the LCD controller has a maximum 'write speed', limiting this.

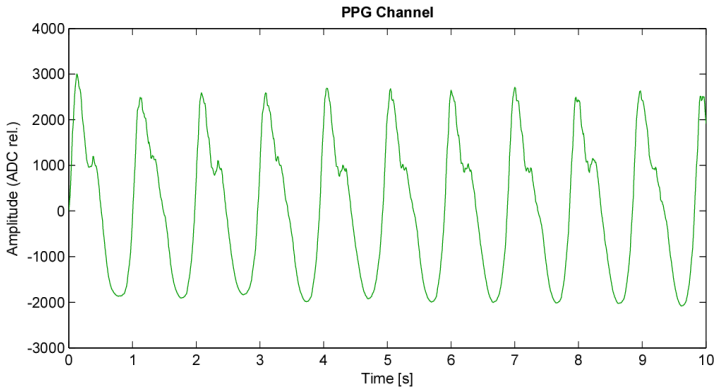

Figure 4.6: *Resulting waveform as captured using the 'black-box' recorder. The result is a faithful representation of a real PPG.*

'virtual ground') then the AC component of a signal will not be proportional to the DC component, as it should be. This is not a large concern for measurements that are taken with a single wavelength where the frequency (used to derive the heart-rate) is the primary focus, but for multi-wavelength measurements that are taken to extract the blood-oxygen saturation (for pulse oximetry, for example), then 'DC-offsets' that are not proportional to the AC component can affect the measurement. Figure 4.7 shows that the envelope of the PPG (an approximation of the AC amplitude) does indeed proportionally match the average light intensity.

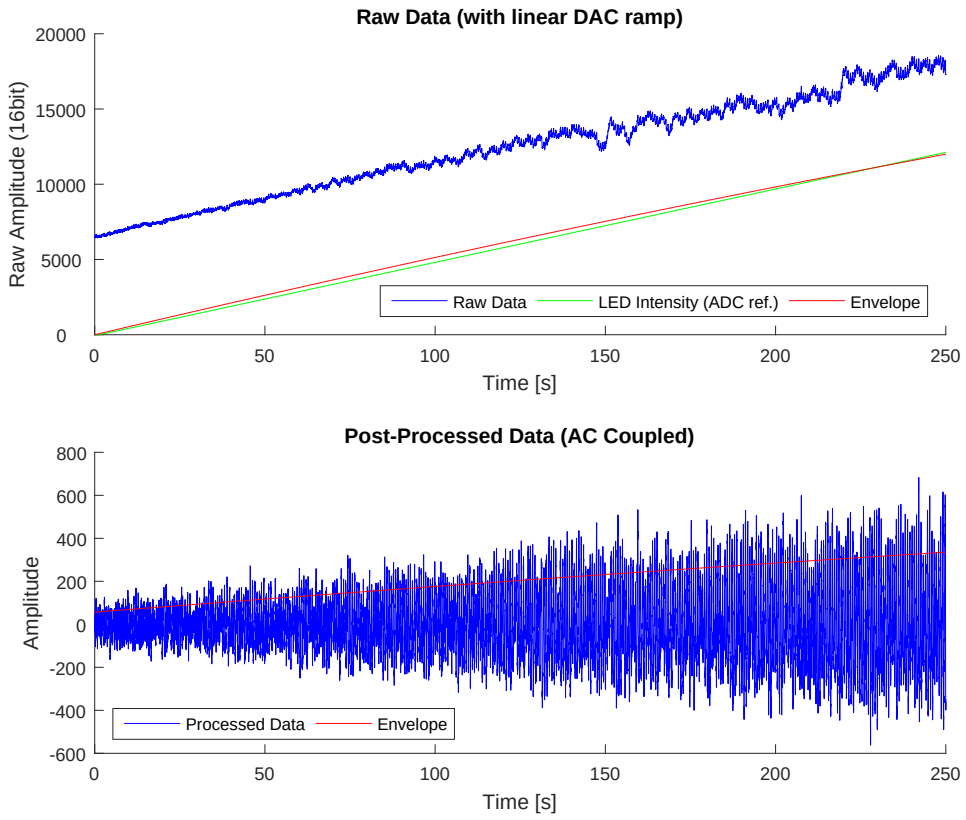

Figure 4.7: *Resulting 'PPG' waveform from the phantom with an increasing LED brightness. The LED brightness has been linearised as previously described. The AC amplitude of the waveform increases approximately equally with the DC component, as expected.*

4.3 Protocol

To collect data for use within this project (mainly in order to test the algorithms as presented in Chapter 4), a participant trial was set-up. The designed and constructed PPG/ECG recorder was used, with two PPG sensors, a 3-electrode ECG set-up, and an accelerometer (as illustrated in Figure 4.8). The two PPG sensors were sewn into a headband such that one was positioned over the centre of the forehead, and one approximately 5 cm to the left side (from the participant's point of view). Each participant performed a routine that started with inactivity (resting for approximately five minutes), followed by a period of activity of their choice (that ranged from walking to running, climbing stairs to jumping, etc), followed by another resting period. The relatively small number of participants recruited

(five, in total, of mixed sex and ethnicity) reflected recruitment problems[4]. The study was approved by the University of Nottingham's Research Ethics Committee (reference number: 2014–140).

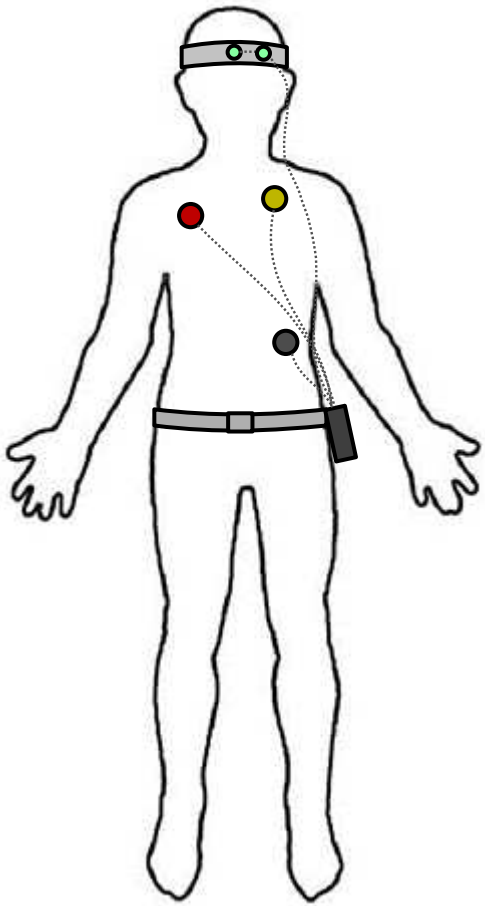

Figure 4.8: *Illustration of the positions of the ECG electrodes and PPG sensors. The accelerometer was integrated into non-centred PPG sensor head.*

As the participants were moving around, away from the lab where the experiments started, the lighting conditions were variable - a mix of ambient day-light and office lighting.

[4]Due to the ethical requirements set out by the University, potential participants could not be approached, they would have to request to take part - few did.

4.3.1 Trial Data

Figure 4.9 shows an overview of the five participants' PPG data, including a 'zoomed in' version that is representative of the signal that exists during the activity period. The walking and resting 'activities' show a strong, artefact-free, signal, for example; whereas the running and 'adjusting sensor' activities show the presence of artefacts and/or larger natural (physiological) variations in the signal.

One common feature of all five participants' data is the general baseline trend (increasing quickly at the start, before slowing, and flattening out). Specifically, for the first couple of minutes of the recording. This is attributed to the skin under the sensors heating up slightly due to the headband and the fact that it is covered[5]. When the skin temperature reaches equilibrium, the baseline change stops.

Taking just one data-set as illustration of the waveforms that were obtained, Figure 4.10 shows an overview (the full recording) of all four channels of data that were stored. The upper and middle graphs shows the physiological data, two PPG channels and the ECG channel (each after being high-pass filtered, with 0.05 Hz cut-off, for clarity), whilst the lower graph shows the accelerometer data from within the recorder box. Participant 4's chosen activity was running (and quickly climbing stairs[6]), and the cadence of their movement is clearly visible in the accelerometer data.

Figure 4.11 shows a subset of the same data. The 100 sps sampling rate is more than sufficient to extract both the PPG and the ECG. If a more detailed analysis was to be undertaken on the ECG waveform, other than simply using it as a pulse reference, then a higher sampling rate may be required. However, the large spike in each pulse (the 'R' of the 'PQRST'

[5]The LED output power is too low to cause detectable local heating (as was measured with a thermometer with the sensor wrapped entirely within the headband); the total output power of the LEDs does not exceed 50 mW.

[6]The participants that elected to do physical exercise were told not to *run* up or down the stairs, for safety reasons.

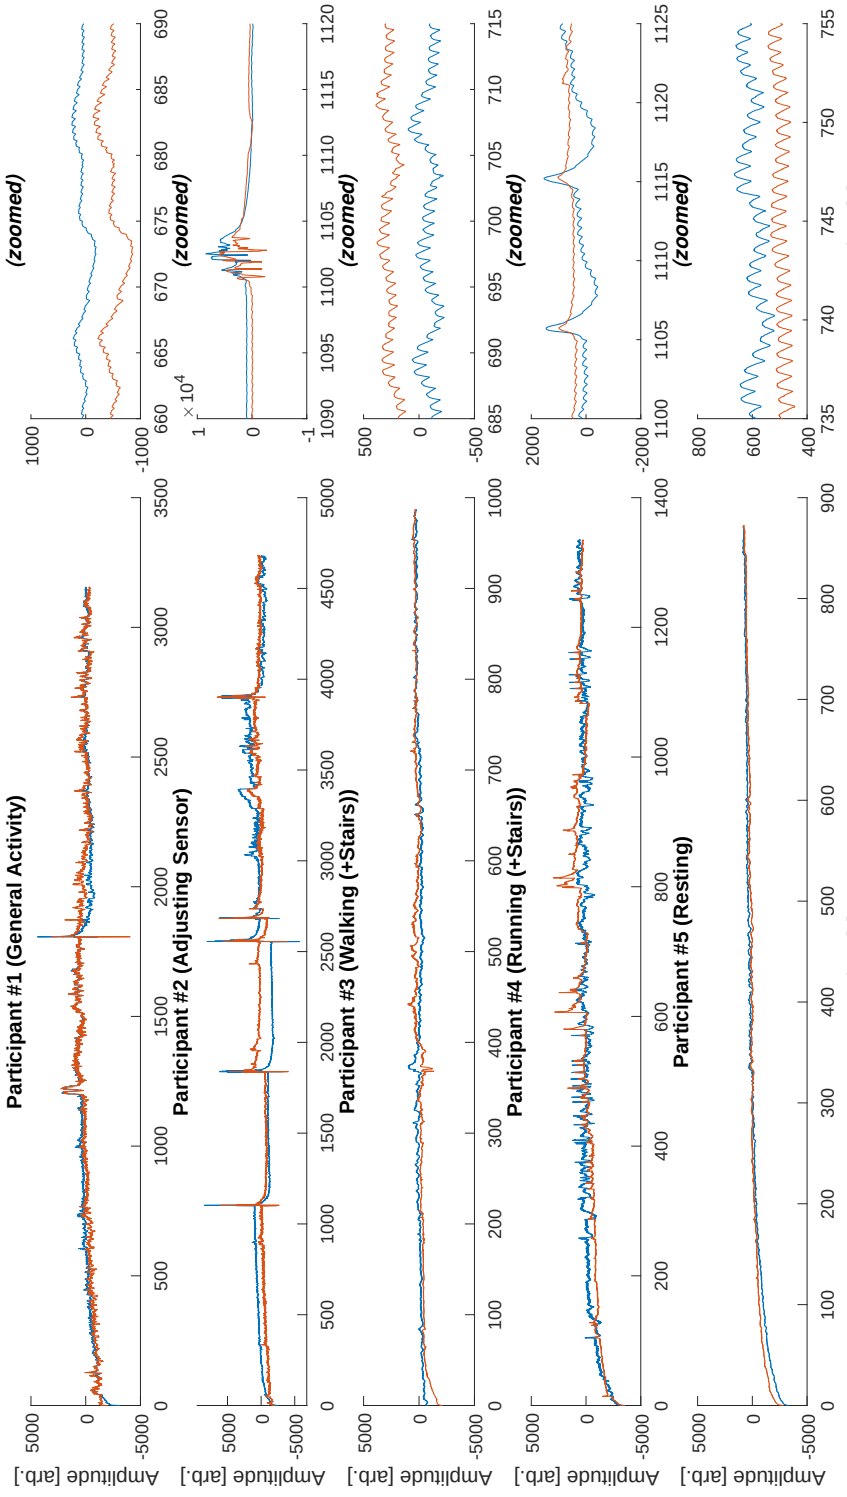

Figure 4.9: *Overview of five participants' data (PPG, two channels). Each participant performed a pre-set activity.*

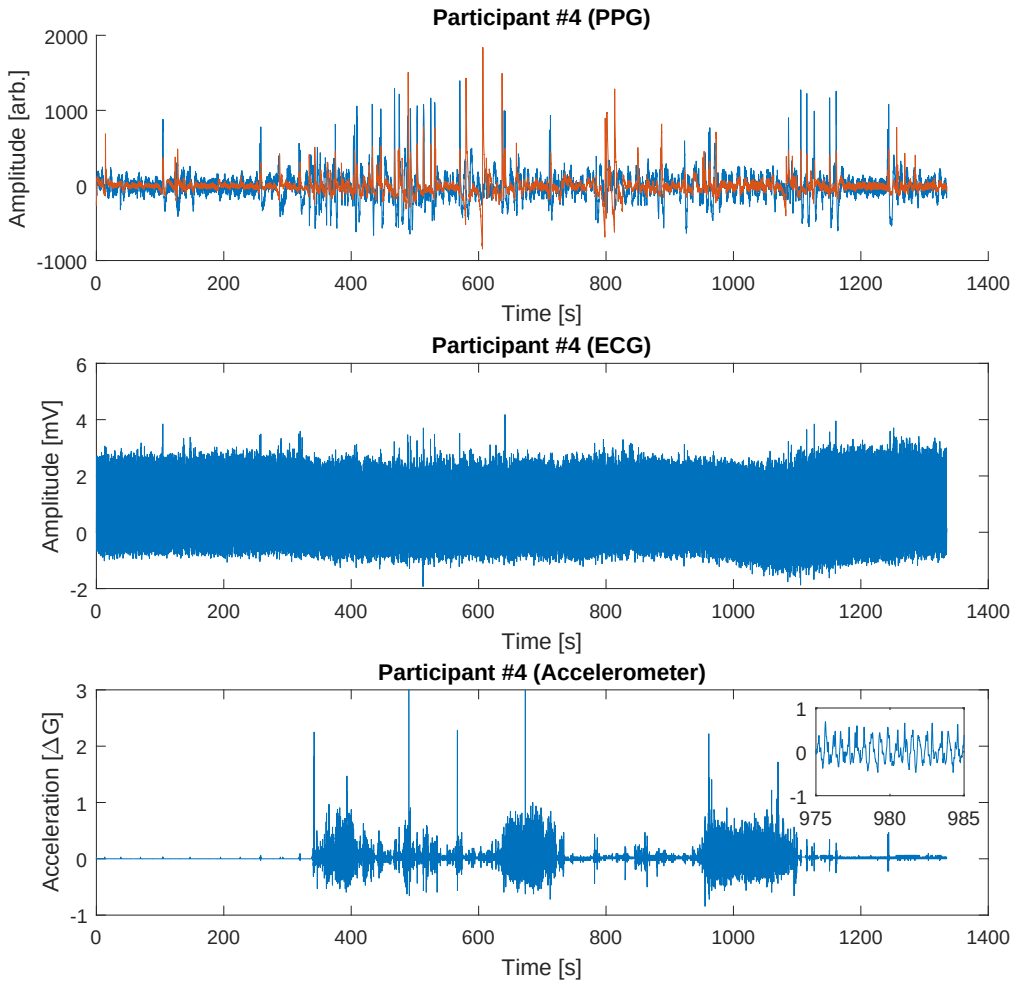

Figure 4.10: *Overview of Participant 4's recording. The duration of this dataset is approximately 22.5 minutes, and the 'activity period' can be clearly seen as large deviations in the lower graph (accelerometer waveform).*

complex), which will be used as the reference location, is still clearly visible.

One risk when extracting PPGs from moving subjects is that, potentially, the movement itself manifests as a component in the PPG signal(s). If the *motion artefact* has a sufficiently large amplitude (compared to that of the PPG), then it may be indistinguishable from a true cardiac-synchronous pulsatile signal, especially if the dominant frequency lies within the heart-rate range. Before any significant processing occurs (as detailed in Chapter 4), a simple visual check was performed on each data-set in the manner depicted in Figure 4.11.

A simple peak-detection algorithm (MATLAB's `findpeaks`) was per-

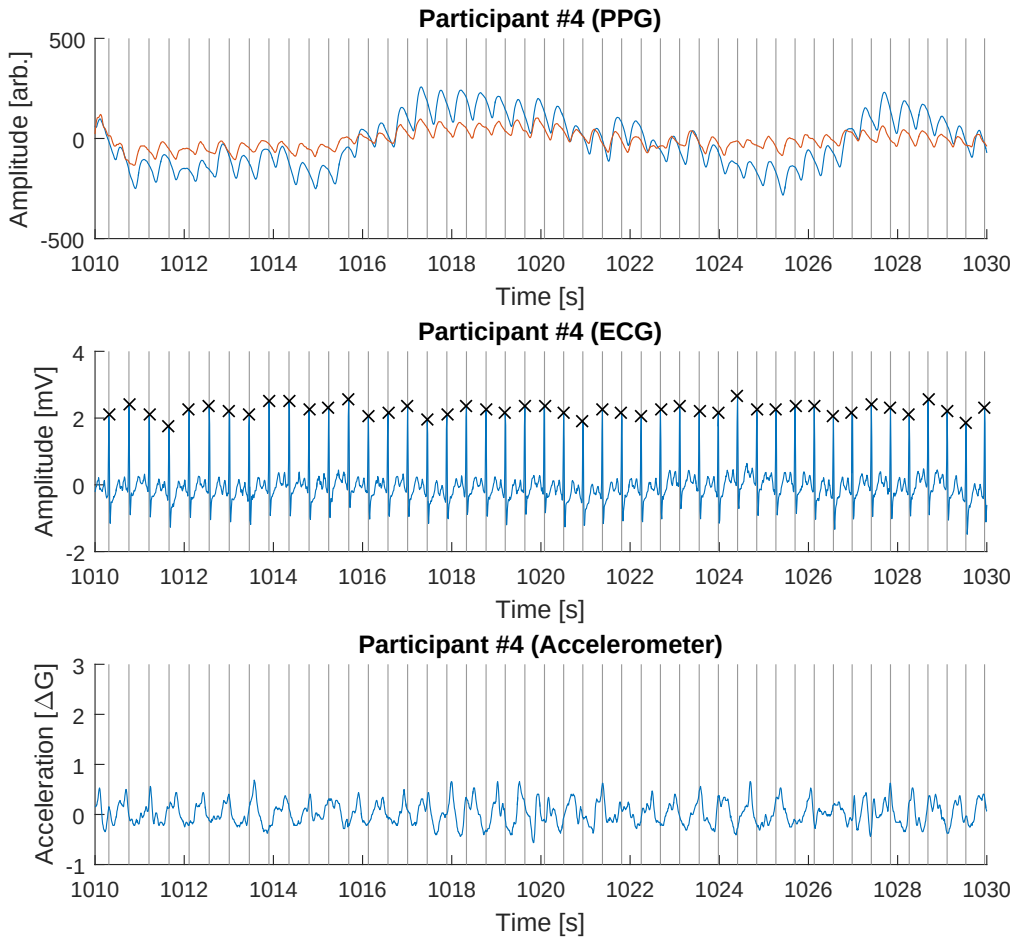

Figure 4.11: *A closer look at Participant 4's data. The grey vertical lines represent the peaks of the ECG pulses, and are duplicated on the other two graphs for reference. The PPG pulses align with these lines, but the accelerometer deviations do not (showing that the pulsatile signal in the upper graph is not caused by movement).*

formed on the ECG waveform, the reference signal, in order to extract the location of the pulses. These detected pulses were then overlaid on **both** the PPG waveforms and the accelerometer signal. The former was to see if the pulses as seen in the ECG matched the pulses observed in the PPG (each cycle of the heart creates exactly one of each pulse, which are aligned, bar a small physiological delay: the pulse-transit-time (PTT)). The latter comparison was to make sure that the repetitive cycles in the accelerometer did *not* match the heart-rate. This was important as it showed that the PPG waveform, that could so easily be formed from movement, was in-fact, not. In all five data-sets, the PPG signals were correlated exactly

with the ECG, showing that the synchronisation worked correctly, and the cyclic movement of the accelerometer was not correlated with the PPG (for all but some severe motion artefacts), showing that the PPG was indeed from the participants.

Now datasets had been collected, the development of the heart-rate and quality-measuring algorithms could commence.

4.4 Quality Measurement

Ideally, the quality level as calculated from the following algorithms should be generalised to allow for a direct comparison between the algorithms themselves. However, as each of the techniques looks at different aspects of the PPG waveforms - some continuous, some discrete - then the outputs would not be easily comparable.

One of the most often used methods of determining whether a PPG signal is "of sufficient quality" is to attempt to extract the heart-rate from it, and then to check whether the heart-rate (either instantaneous or average) is within an acceptable 'pre-set' range [84–87]. Whilst this method provides a quick and simple way of detecting where bad PPG regions are likely to be, it is critical that it is not used on its own as large periodic artefacts that have similar properties to the PPG can be erroneously picked-up as a valid PPG signal. In other words, the heart-rate range algorithms can reliably reject bad data, but cannot reliably accept other data as good (it will generate *false-positives*).

Whilst PPG peak-to-peak (or trough-to-trough) times will ideally yield the more desirable 'instantaneous heart-rate', methods that rely on sectioning the data (FFTs, for example) may produce a more consistent and robust heart-rate output - at the expense of the poorer response times. As detecting the quality of a PPG signal by itself does not yield any use-

ful clinical information (the desired output is most often the heart-rate), methods of extracting the heart-rate, as a means of measuring the quality, are covered first.

4.5 Heart-Rate Detection

Three methods of heart-rate detection are considered now. As stated before, each method will allow for the rejection of poor quality PPG regions, but cannot reliably detect good regions, rather, additional techniques will be required, and are considered afterwards.

4.5.1 Peak & Trough Detection (Time Domain)

A single 'pulse' in an electrocardiogram (ECG) has a very distinct shape that includes a single high amplitude 'spike' (the 'R' of the "PQRST" complex); peak detection algorithms can extract the position of this peak when no artefact is present due to its short duration and dissimilarity to other components within the same pulse. When the more rounded and less distinct photoplethysmogram pulse in used in its place, the algorithms required to extract the pulses (peaks and troughs), by necessity, become more complex. Subtleties in the shape of the PPG, particularly when the dicrotic notch is included, cause problems when a reliable algorithm is required. In this section, the algorithm used to extract the ECG pulses in the datasets is not discussed, focussing instead on the PPG.

As with most algorithms, several stages are required to calculate the final outcome. Each stage usually refines the outputs from the previous stages, removing errors, and correcting mistakes. The main stages of the algorithm designed and used are discussed now.

4.5.1.1 Filtering

As this particular technique does not use the low frequency baseline of the PPG, it can be filtered out (as well as the much higher frequencies caused by flickering lights, etc). The filter as described in Section 2.6.3 is optimal in this case, where only components within the heart-rate range are required. A Butterworth bandpass filter is implemented with a pass band of 0.66 Hz to 9.00 Hz. This has an additional advantage in that the mean is removed so scaling does not affect the offset.

4.5.1.2 Naïve Peak Search

Here an 'overly enthusiastic' peak (and trough) detector is used. This detector detects all significant[7] extremes, treating peaks and troughs separately. For PPG signals, the dicrotic notch is often picked up as a separate extreme (whether it is detected as a peak or trough depends on the polarity of the signal, or course), and can be attributed to sudden changes in calculated heart-rates.

The algorithm used for this search is the `findpeaks(...)` function from MathWorks® MATLAB, with the `'MinPeakDistance'` argument set to `fix(0.33 * sps)`, such that peaks that occur more often than 3 Hz are not picked up (see Figure 4.12). This is not a perfect solution as PPGs of people with lower heart-rates will be more likely to have secondary peaks picked up (such as the dicrotic notch, or artefacts). Ideally, the `'MinPeakDistance'` argument would be set to a function of the current heart-rate; however, if the heart-rate was known, then this technique would not be needed in the first place.

[7] "Significant" in this case refers to extrema that are a minimum distance (in time) apart. The distance is set such that the maximum heart-rate can still be detected.

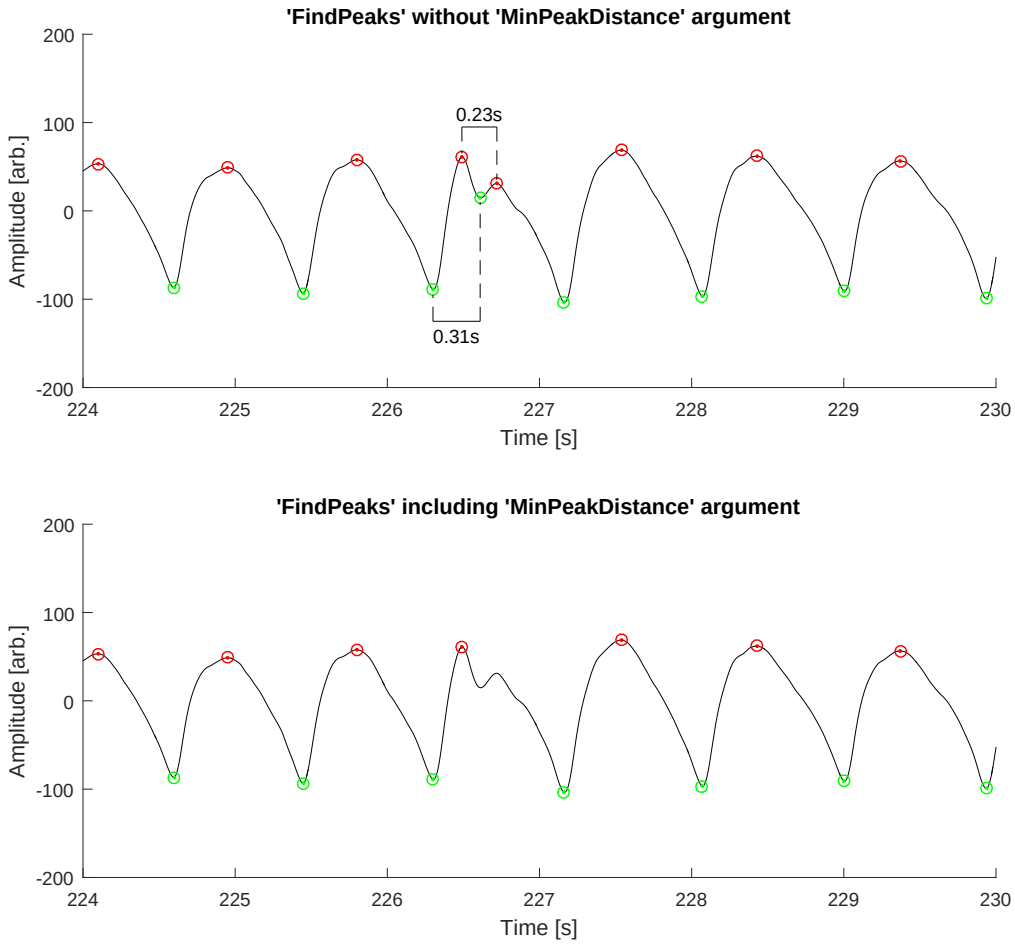

Figure 4.12: *Illustration of how* `findpeaks`*'s* `MinPeakDistance` *remove peaks that are too close to each other.*

As the `findpeaks(...)` function only look for 'peaks' (positive extrema), then the function was also called with the inverted PPG to find the troughs. The values (amplitudes) of the detected 'troughs' were then inverted after the function to return them to their original polarity (see Listings 4.1).

```
[ppg1_pv, ppg1_pi] = findpeaks(+ppg1_filt, ...
      'MinPeakDistance', fix(0.33 * sps));
[ppg1_tv, ppg1_ti] = findpeaks(-ppg1_filt, ... % Negate PPG.
      'MinPeakDistance', fix(0.33 * sps));
ppg1_tv = -ppg1_tv;   % Negate trough values.
```

Listing 4.1: *Extract from heart-rate finding algorithm.*

4.5.1.3 Alternating Pattern

To prevent spurious heart-rate fluctuations due to extra pulses, the results from the peaks and troughs are looked at separately, and a simple rule is used to remove some of the extra peaks. This rule removes duplicate extremes that do not fit into an alternating pattern. The amplitude of the PPG alternates high-low-high-low etc., with each pulse being aligned within a low-high-low (trough-peak-trough) triplet, i.e. each pulse shares the same extremes with its neighbours. If two of one extreme fall within two of the other extreme then one is removed; specifically, the one with the lower magnitude relative to the expected extreme[8]. This is run on both the peaks and the troughs separately to ensure an alternating pattern all the way through the recording.

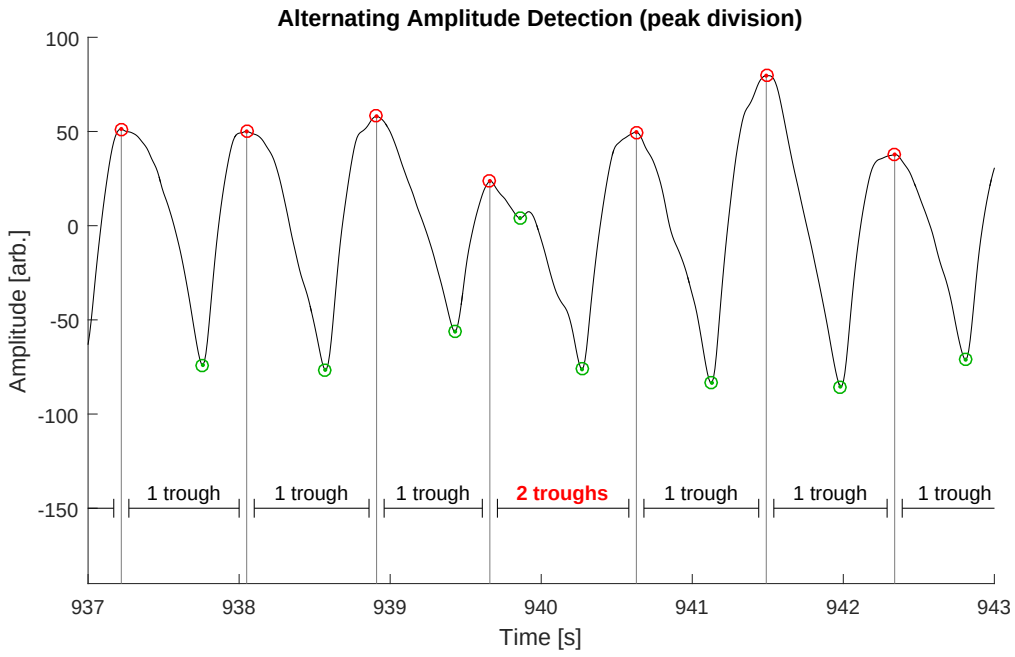

Figure 4.13: *Illustration of how the check for alternating extremes is performed. This example only shows the first pass where the number of troughs are counted between each pair of peaks. The same routine is run afterwards were peaks are counted between pairs of troughs.*

[8]Larger amplitudes (more positive) for troughs, smaller amplitudes (more negative) for peaks

In Figure 4.13, there are two troughs found between the two peaks straddling 940 seconds, of these two, the lower one (the right-most) is kept, and the other is discarded.

4.5.1.4 Overlapping Envelopes

The final step looks at the envelopes that are created when each of the peaks and troughs are interpolated. For an ideal PPG signal that has no artefacts, the two envelopes should be perfectly parallel, and would therefore not cross. When artefacts are present that cause local amplitude fluctuations, the envelopes may cross where additional extrema exist. Errors like these may be caused by an exaggerated dicrotic notch, or higher frequency fluctuations (artefacts). Figure 4.14 shows one such example before the correction takes place.

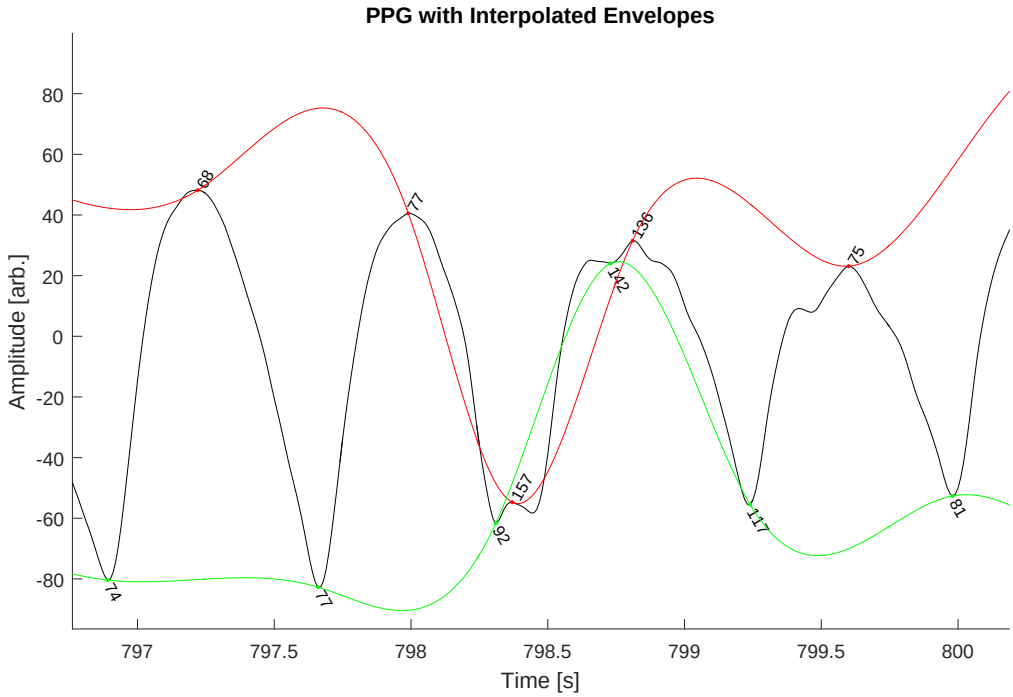

Figure 4.14: *An example of where the amplitudes of the peaks and troughs can erroneously swap. When this occurs the envelopes will cross.*

If, of the two local minima at 798.5 seconds, the right-most had been lower, and therefore chosen as the trough, the previous stage (that looks for alternating patterns) would have discarded the falsely detected maxima

(peak). Since this was not the case, there exists a correctly alternating pattern that consists of an extra 'pair' of extremes. This extra pair will be formed when small fluctuations exist in two neighbouring pulses as seen in Figure 4.14, and will have 'opposite' amplitudes - the peak will be lower than the trough. Unfortunately, it is not quite as simple as removing peaks that are 'too low' or troughs that are 'too high' (see Section 4.5.1.6).

4.5.1.5 Peaks & Troughs Results

In order to compare the results of the algorithm, the heart-rates as calculated from the ECG and PPG waveforms are overlaid on top of each other to allow for quick visual inspection. Three small sections are illustrated in Figures 4.15, 4.16 and 4.17 showing the effectiveness of the algorithm and its small susceptibility to artefacts.

When no artefacts are present in a signal, such as illustrated in Figure 4.15, the detected peaks and troughs individually produce instantaneous heart-rates that are well within 5 bpm of the reference[9]. Under these conditions it may be possible to average the two heart-rate outputs to obtain a more accurate result; however, the hearts-rates from the peaks and troughs are being kept separate due to reasons that will become clear.

Since each pulse contains a single dicrotic notch (that may not necessarily be visible in the waveform) that can interfere with the peak detection algorithm (such that there are two possible peaks that could be chosen), local perturbations may exist in one of the results; specifically, the heart-rate calculation that uses the 'peaks' (when the signal represents attenuation). Small fluctuations are visible in Figures 4.16 and 4.17 (red lines). It is clear how these exist by simply observing the switching of peak position at around 350 seconds in Figure 4.16; the correct peak is detected first (the first peak before 350 seconds), followed by an incorrect peak (the dicrotic

[9]The range of ±5 bpm was chosen arbitrarily as a way of illustrating the similarity between the values.

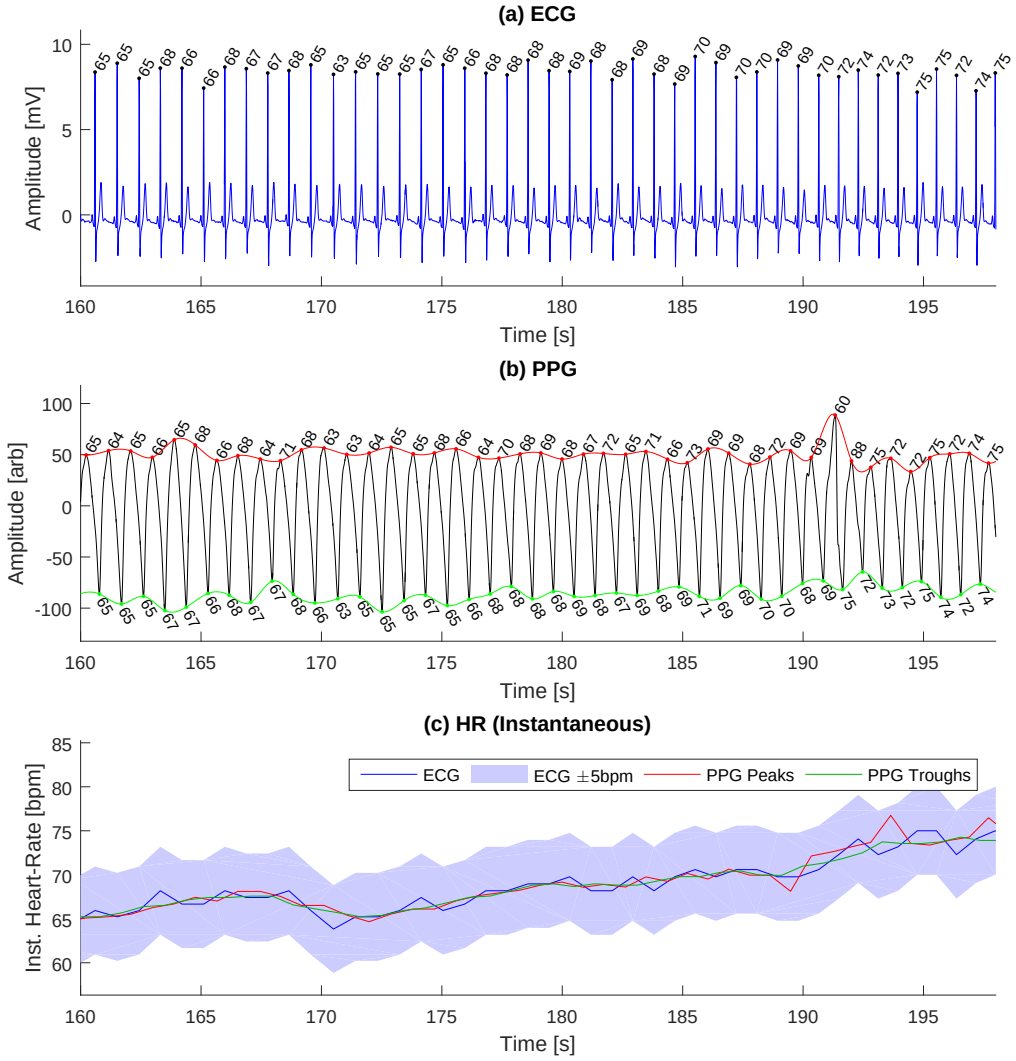

Figure 4.15: *Example of instantaneous heart-rate detection on a 'clean' (artefact-free) data set. Values at each peak (of the ECG and PPG), and trough (of the PPG) represent the instantaneous heart-rate for the pulse.*

notch is selected), followed by the correct peak again.

If the 'wrong' extrema is chosen then the instantaneous heart-rate will suddenly either increase or decrease. However, as this extrema is also used in the instantaneous heart-rate calculation for the next pulse, the following result will also have an error, just with the opposite 'direction'. In other words, any error caused by the dicrotic notch will 'average out' in the next pulses (but there will always be an error in at least two values).

At the expense of a small delay (of one or two pulses), a simple averaging filter that averages consecutive heart-rate readings will mostly correct for this anomaly. This can be clearly seen in Figure 4.18.

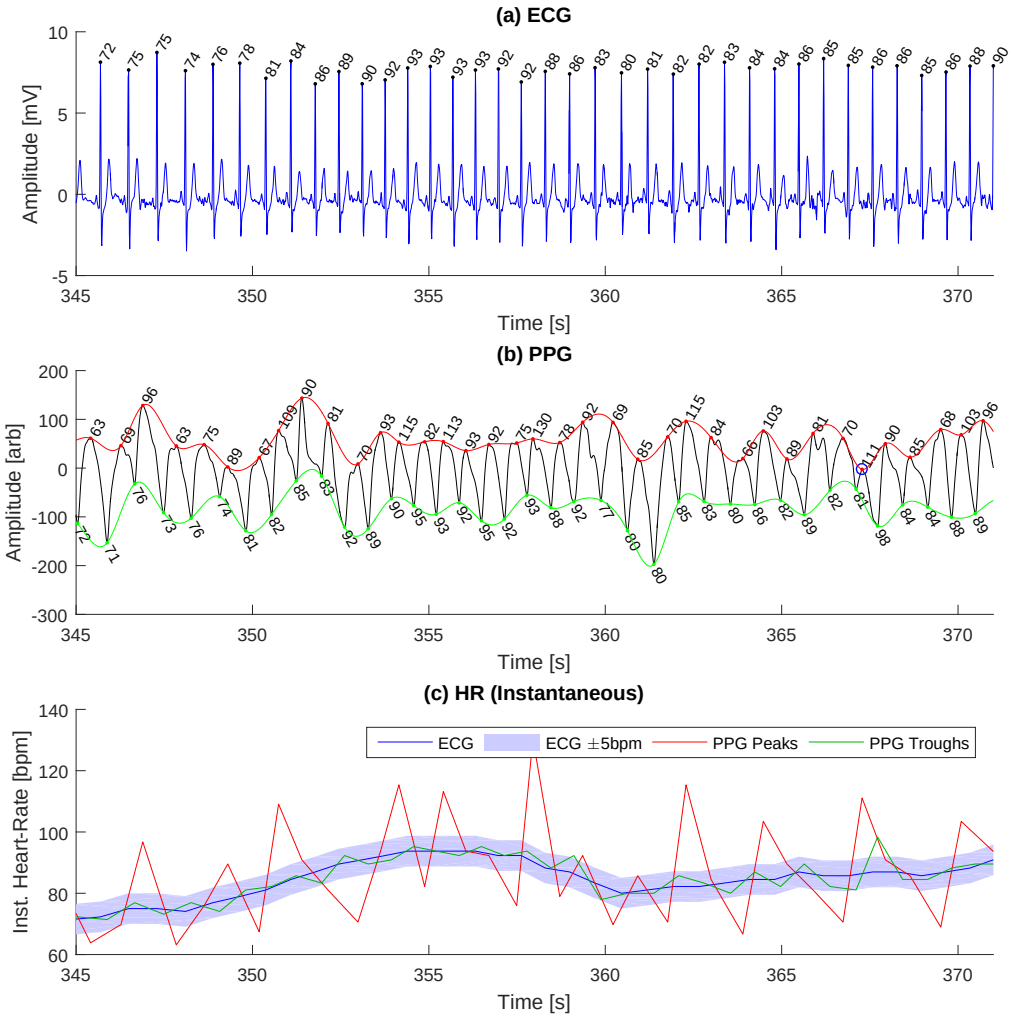

Figure 4.16: *Example of instantaneous heart-rate detection on a data set which has a small amount of motion artefact present, causing a minimal amount of disruption to the heart-rate measurement (the 'trough' values are still sufficient to correctly determine the heart-rate).*

4.5.1.6 Peak-Trough Correction Methodology

The technique presented contains no so-called "magic-numbers", i.e. constants that are specifically 'tuned' to the dataset being used. In other words, there are no numbers that do not have any explicit meaning within the calculations being performed. In fact, the *only* time any constants are defined in the code are when the initial filter is concerned (regarding the pass-band frequencies). This means that any dataset can be used without having to 'tweak' settings or pre-calculate statistics. Additionally, as all functions and techniques used in the algorithm only use the current and

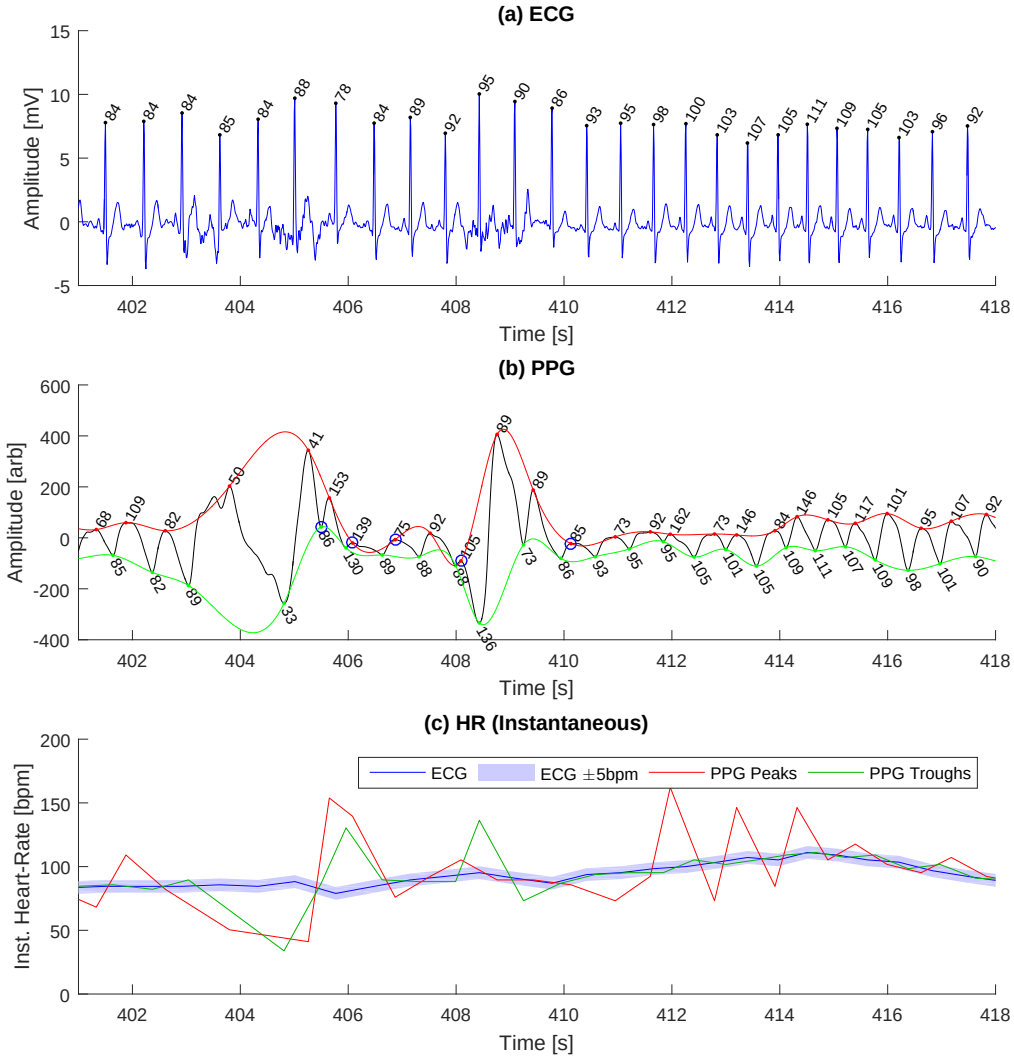

Figure 4.17: *Example of instantaneous heart-rate detection on a data set which has a large (albeit short-lived) artefact, causing significant disruption of the detected heart-rate.*

previous data points (there is no forward scanning of data), the algorithm would be easily converted into one that accepts streaming data - i.e: can be used with live data with virtually no delay.

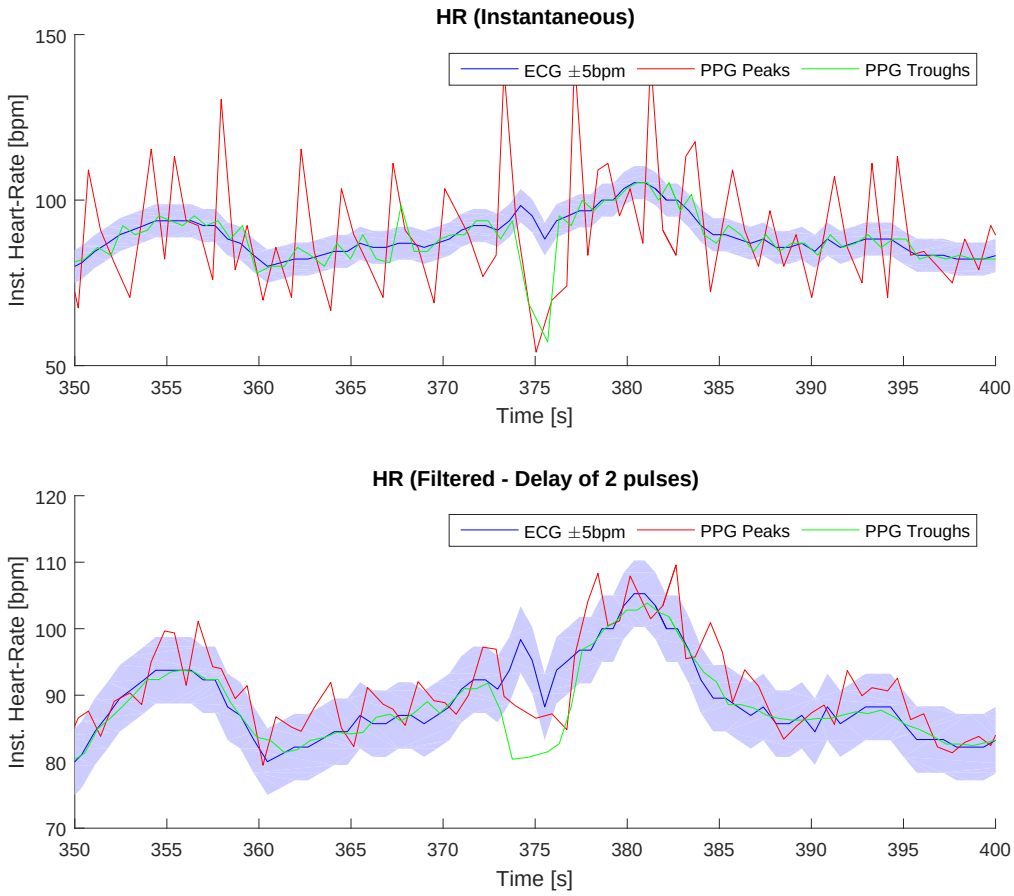

Figure 4.18: *Effect of smoothing the heart-rate values. Note that the heart-rate 'line' created from the troughs remains mostly unchanged as this did not have the same fluctuating characteristic as the other line.*

4.5.2 Frequency Analysis

As previously stated, when performing frequency analysis on a data-set, each frequency spectrum must originate from a window (in time) of the data; i.e. you cannot estimate a frequency from a single point of data. Extracting heart-rates from PPG data using transformations into the frequency domain must therefore induce a 'lag' in the output. In many applications such a delay would be acceptable, but in some, it is not. The delay is exactly equal to the window size of the transform - a five second window will result in a five second 'lag', for example (ignoring processing times). The term 'window' in this context is referring to the sub-set of data that is currently being processed; to generate the result from the whole set of data, the calculations are repeated with a new sub-set (of the same size),

but with a different 'centre'. When the amount the 'window' moves is less than its width, some data is 'reused', for example.

When transformation techniques such as Fourier Transforms are used (either Discrete (DFT), or Fast (FFT)), the 'window' size, that equates to the length of the data being processed, also defines the frequency resolution of the output. The more data that exists, the more accurate the end-frequency 'bins' are. This leaves a problem of compromise: high frequency resolution versus small temporal lag.

Initial bounds can be set by considering the expected heart-rate range. If multiple pulses are required to get an accurate response, then the reciprocal of the lowest heart-rate would be the lower bound. If 40 bpm is taken to be the lowest heart-rate, then $samplerate * 60/40$ would be the minimum number of samples to use. At 100 samples per second (sps = f_s), the minimum window size would be 150 samples (= N), equating to 1.5 seconds. The "Nyquist Frequency" of 50 Hz is the maximum resolvable frequency that is achieved at the sample-rate [88], and the resulting frequency resolution (the difference in frequency between successive 'bins') is $f_s/N = 100/150 = 0.6\dot{6}$ Hz = 40 bpm. It is quite clear that this is nowhere near good enough for heart-rate extraction.

> The Nyquist-Shannon sampling theorem is not being used 'correctly' here as there are more constraints that must be placed on the signal before the $f_s/2$ condition is true (the signal must be band-limited, for example). However, as this particular example is being used to give an *idea* of the values involved, and how the FFT method is affected by them, this condition is ignored for clarity. In a real-world scenario, the sampling rate could be made to be much higher, if a method like this were to be used.

Flipping the problem on its head: what would the required window size be in order to achieve a frequency resolution of 1 bpm? The equivalent heart-rate (resolution) as a frequency would be 1/60 Hz = $0.016\dot{6}$ Hz. At the same 100 sps as before, the number of samples required would be f_s/f_r = 100/0.0166 = 6000 samples yielding a delay of 60 seconds. This unreasonable result for the preferred resolution would create many problems in many acquisition systems. It is simply too high a delay.

In addition to the larger delay, a larger window size presents another, different, problem. For applications where the immediate heart rate is critical (medical applications), as opposed to where it is simply for casual information (fitness), the local changes in heart-rate - its modulation (heart rate variability (HRV)) - may contain important information. An FFT-based method with a window size of 10 seconds will not 'see' a missing beat, for example, as the resulting 'peak' frequency component (the FFT 'bin' with the highest amplitude) will not contain any information about *when* the peaks occurred[10]; the missing pulse will be lost. This also suggests that shorter window lengths would be preferable, so as to be useful for applications where information about the pulses themselves (and whether they are present or not) are required.

By reducing the window size to approximately 5 seconds[11], the frequency resolution is now f_s/N = 100/512 = 0.1953... Hz = 11.72 bpm (see Figure 4.19, red line). Although this window size is a fair compromise between frequency resolution and temporal lag, it is still not particularly useful for many applications. It can, however, be improved with a small amount of additional processing.

Each bin in an DFT/FFT output represents a single frequency. If one particular bin were to represent, say, 100 Hz, then a sinusoid with the same

[10]The spectra that results will look similar (having the same 'peak' frequency), but will be 'spread out' as the signal would no longer be periodic.

[11]The DFT method used (from MATLAB) requires a 'power-of-two' sized data-set. The actual number of samples chosen for the window is 512, yielding 5.12 seconds.

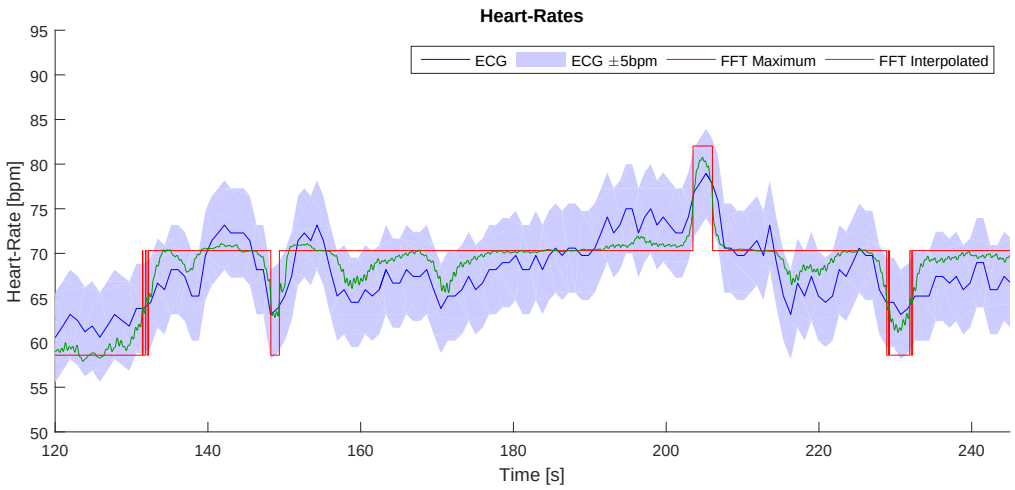

Figure 4.19: *Section of the results from a single data-set. The blue line represents the 'actual' heart-rate as calculated from the ECG (peak-to-peak). The red line represents the peak frequency from the FFT of the PPG data. The green line represents the interpolated frequency from the FFT (explained in text).* ***Note:*** *Due to the flexibility and non-time-critical nature of post-processing, the heart-rate results are shifted back to correctly align with the other data. In other words, the temporal lag has been artificially removed. This is not possible to achieve with 'live' processing and is simply implemented here to improve clarity.*

frequency would appear entirely within that bin. If the next bin 'up' (the next higher frequency) were to represent 101 Hz, then the same frequency would, likewise, entirely reside in that bin. As a discrete (sampled) input to an FFT produces a discrete number of output bins, in this example, there would not exist a bin at 100.5 Hz. If such a frequency were to be input to the function, then the resulting value in the frequency domain would be 'spread out' around the neighbouring bins (see Figure 4.20).

The fundamental frequency of a PPG, the source of the heart-rate, will vary along a non-discrete domain; therefore, the frequency components resulting from the FFT will, although always be spread out, have their dominant peak 'snap' to the nearest bin (the bin with the largest magnitude). However, for a given input frequency, there will be a *bias* away from the dominant bin (for frequencies that do not lie exactly on one). It may be possible to improve this interpolation if a different technique is used. Figure 4.21 shows three spectra for sinusoids with slightly different

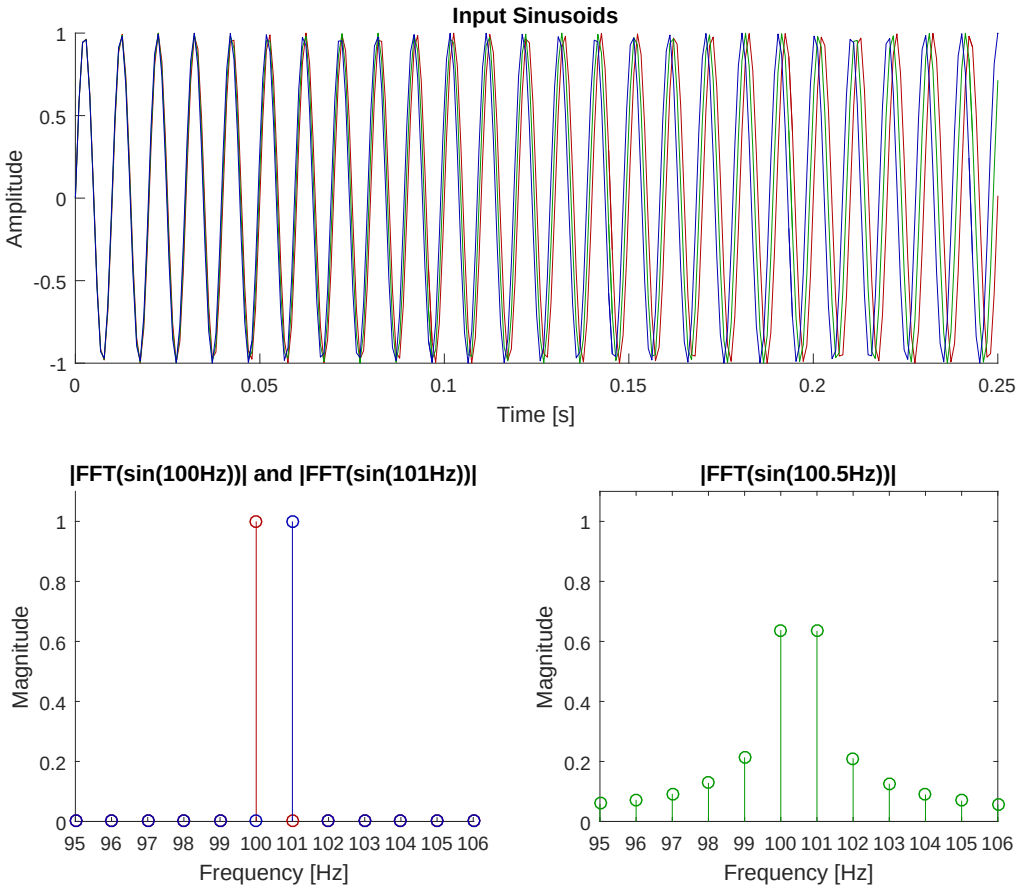

Figure 4.20: *Example of FFT 'spread', whereby frequency components that do not lie exactly on frequency bins present themselves incorrectly as a range of frequencies centred at the correct virtual position. Vertical axis represents the magnitude of the FFT.*

frequencies (100.4 Hz, 100.5 Hz and 100.6 Hz). The 'bias' is clear to see, especially when a polynomial is fitted around the top-most values. In fact, by taking the interpolated maximum of the fitted curve, a *more* accurate frequency can be obtained.

It must be noted, however, that the 'biassed' peak is not a perfect representation of the actual frequency. The green line in Figure 4.19 shows the interpolated maximum in direct comparison with the standard 'max' function. Whenever there is a jump between bins in the standard approach (red line), the interpolated value also jumps; however, in the periods where the standard approach gives a flat response (between 160 and 180 seconds, for example), the interpolated maximum does a better job at matching the reference output (blue line). In this particular part of the data set,

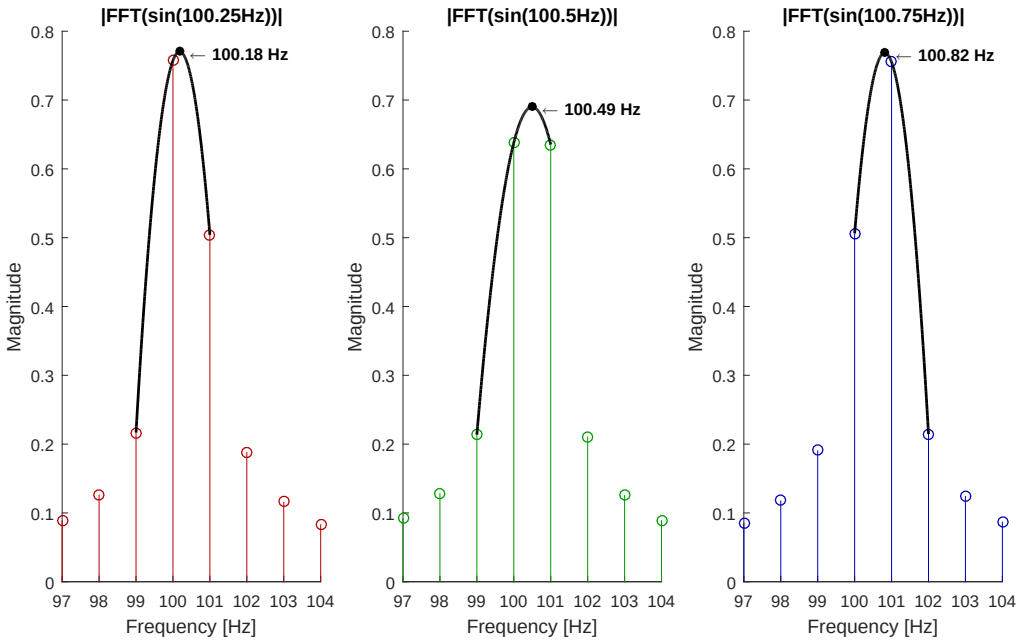

Figure 4.21: *FFTs of three sinusoids whose frequencies are 'off-centre'. By fitting a polynomial or spline to the peak and surrounding two points, an approximation to the 'real' maximum can be gained.*

the interpolated line lies entirely within the ±5 bpm band of the reference signal. Additionally, this technique will work less well when the frequency components of the signal are many, and have fewer defined 'peaks'.

4.5.2.1 Error Detection and Correction

When an artefact occurs in the PPG signal, the detected frequency of the heart-rate may jump to a different frequency (typically one that is much lower, due to the often non-repetitive nature of most artefacts). For a simple algorithm, such as the one developed here where no 'tracking' of the heart-rate takes place, it is simpler to ignore regions where an expected artefact takes place. This rejection can be achieved by looking at the first derivative of the detected heart-rate.

Whenever a sudden change in heart-rate occurs, such as illustrated in Figure 4.22, it is often (but not always - see later) paired with another sudden change a short while later, after the artefact has ended. In the datasets that were used for these calculations, artefacts' frequencies have

not been observed to be higher than that of the PPG **and** be of a higher amplitude. The amplitude, in fact, for observed artefacts have tended to be inversely proportional to their frequency. This is the reason for each artefact causing a drop in the frequency of the highest-amplitude peak in the FFTs.

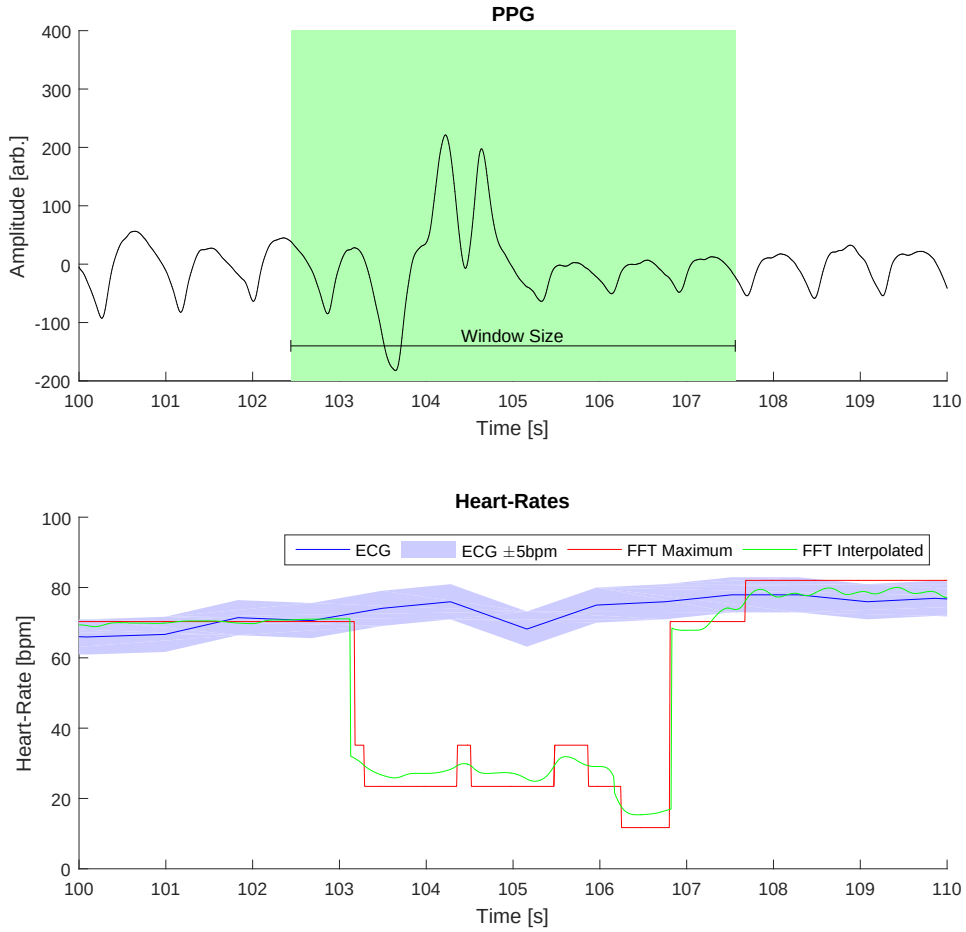

Figure 4.22: *Example of a sudden 'step' in the detected heart-rate due to an artefact.*

Figure 4.23 shows two frequency spectra for regions with and without artefacts (the window for the signal with the artefact is the same as the one in Figure 4.22). Regions that are bounded by a sudden (within 10 ms) change in detected frequency (of more than 10 bpm), are considered 'bad'. These regions are replaced by a heart-rate 'guess' based on the last and next good signals. Effectively, this guess consists of a straight line that bridges the gap.

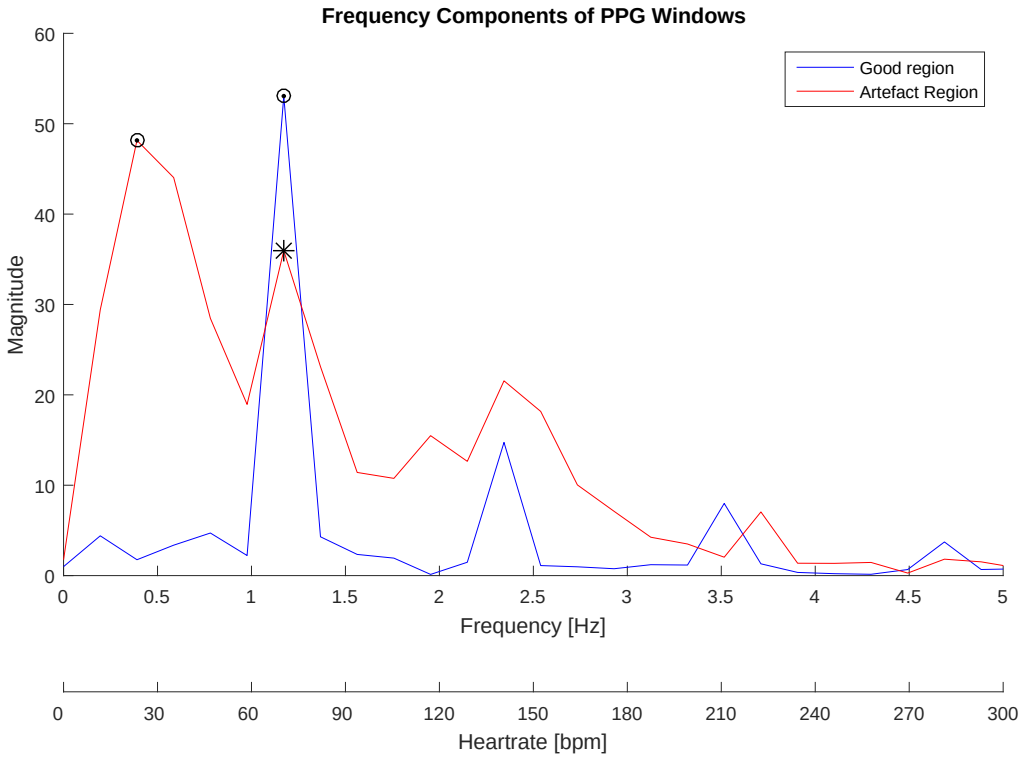

Figure 4.23: *Illustration of why an artefact causes a sudden drop in detected heart-rate. As artefacts of lower frequencies tend to have higher amplitudes than higher frequencies, at some point as the artefact 'grows', the peak frequency will suddenly snap to the artefact. The circled peaks represent the detected maximum amplitude frequency, the starred peak represents the frequency associated with the heart-rate.*

4.5.2.2 Frequency Analysis Results

For regions of data which are free from artefacts (such as illustrated in Figure 4.24), or have artefacts that have a lower amplitude than that of the PPG, the frequency analysis method works reasonably well. Despite the temporal lag in the output, and with the additional processing to "improve" the resolution, the heart-rate that is detected is within ±5 bpm of the reference.

When larger variations in the PPG signal exist, either as a result of artefacts, or simply by large baseline drifts as is shown in Figure 4.25, the frequency analysis will briefly show an incorrect result. In this particular case, no 'artefact' exists in the traditional sense; instead, due to small arrhythmias (bradycardia events, which are clearly visible in the ECG trace),

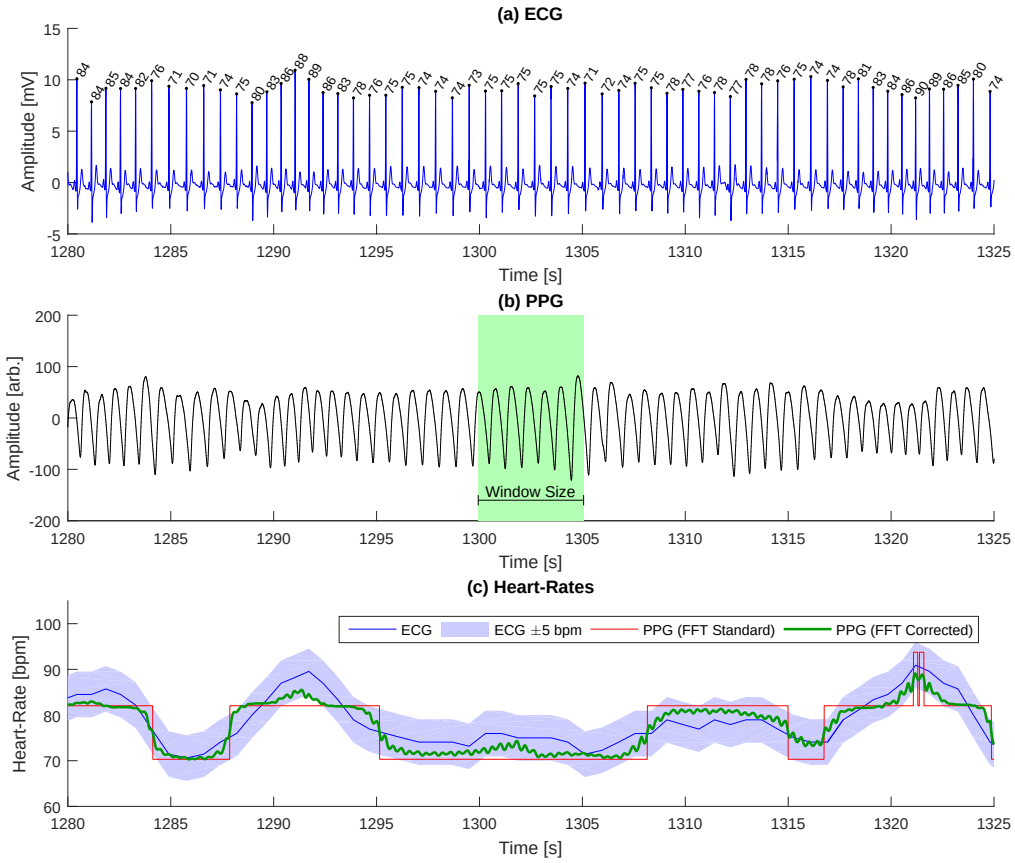

Figure 4.24: *A good section of PPG showing the resulting heart-rate as calculated via the FFT/DFT method. A 5.12 second rolling window was used.*

the baseline of the PPG shifts dramatically causing the 'low frequencies' that are dominating. This shift may be attributed to changes in blood pressure due to other events that are not directly related to the heart, similar to that which is caused by forced or restricted breathing [89] (heavt breathing due to the activities may be the case of this). The 5-second window is very clearly seen here - it is the duration of the 'errors' caused by each of the two arrhythmias.

The temptation would be to filter the signal with narrower pass-bands in order to remove the turbulent baseline drift, however, as explained in Section 2.6.2, 'over-filtering' the data could cause real artefacts to be considered as pulses. The filters used to pre-condition the PPG have been chosen so as to not corrupt the data in this way.

Due to the 5 second window and the additional 'error correction', these

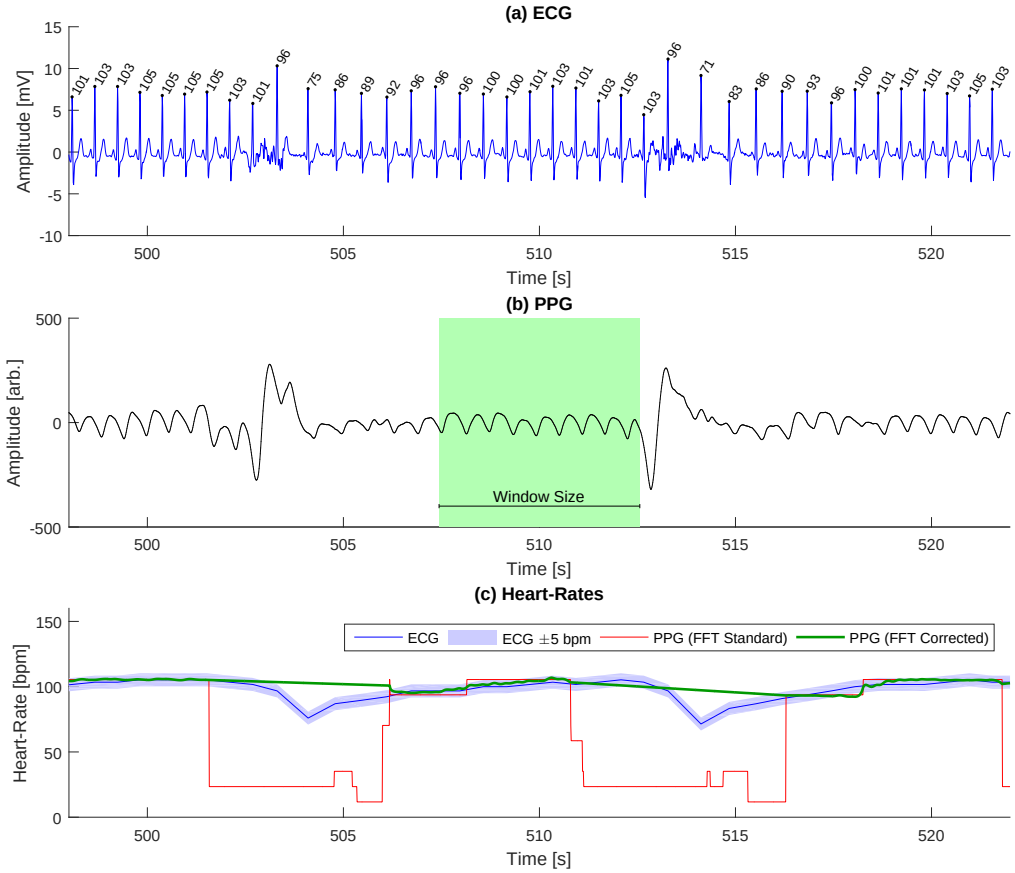

Figure 4.25: *The two small fluctuations in the PPG in this figure are not necessarily 'artefacts' as they are part of the signal: the heart-rate has instantaneously changed (as can be seen in the ECG). The lower frequency components in the signal have dominated and have caused a sudden incorrect evaluation of the detected heart-rate.*

two sections have been 'ignored' and the previous valid heart-rate has been linearly interpolated to the following valid heart-rate effectively bridging the gap. For artefacts and irregularities in the signal, this is a reasonable approach to 'filling in the data'.

As is expected with biological systems, there is always a situation that falls outside of the norm. Between 450 and 455 seconds in Figure 4.26, there is a sudden drop in heart-rate that is not caused by the same mechanism as the previously mentioned arrhythmias. In this example, the PPG baseline does not change. Due to the lower frequency resolution of the FFT based method, the heart-rate drop is considered by the algorithm as an error. It has, in fact, just cleared the threshold, causing it to be flagged. As

this is not an artefact, there is no expectation for the heart-rate to 'jump' back up to previous value. This causes a long region that is declared as erroneous (the artefact region is actually terminated by one of the following arrhythmias), and is longer than desired.

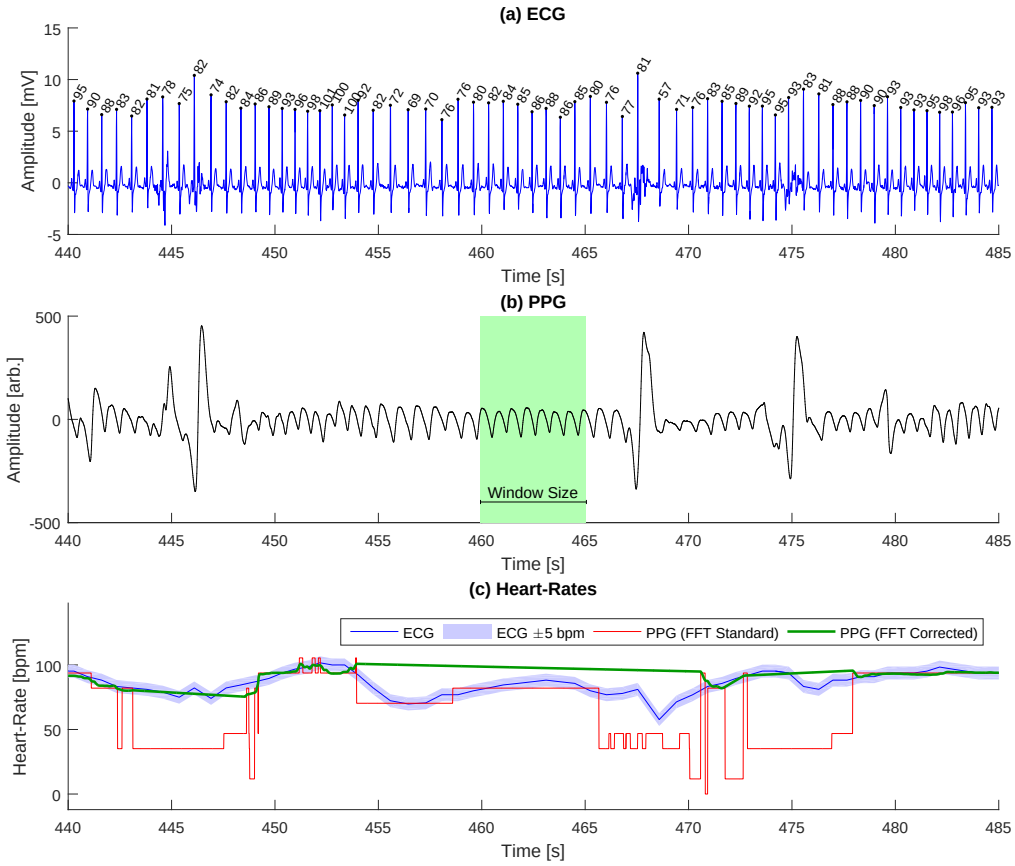

Figure 4.26: *Due to the windowing technique that is necessary for the FFT method, there are instances where a perfectly good PPG signal is marked as bad. In this particular case, a sudden (valid) decrease in heart-rate causes the error detection part of the algorithm to believe there is a bad section.*

During regions of the PPG with no artefacts, and those with moderate 'non-repetitive' artefacts, the frequency analysis method is able to obtain a heart-rate that is within five beats-per-minute of the reference. There are some situations that cause the basic algorithm to fail however, particularly when the artefact has a fast onset, causing an easily detectable drop in heart-rate, but then having a slow recovery. One such example is presented in Figure 4.26; at 454 seconds, a drop in heart-rate exists that is

incorrectly detected as the start of an artefact, however, the small 'recovery' at 459 seconds is not sufficient for the algorithm to register the end of the 'artefact'. Instead, it uses the following large positive 'heart-rate change' resulting in the large 'bridge' that spans the artefact-free region. This method, in the same manner as the peak/trough technique, was designed for use with streaming data. Under streaming conditions, it is not possible to 'look-ahead' at future data as it does not exist at the time. Short of delaying the output further (on top of the delay caused by the windowing itself), it is not easy - and potentially impossible - to predict when the 'end' of the artefact will be. The algorithm must therefore wait until a significant event occurs before resetting.

Alternatively, as a way of preventing the heart-rate 'drops' in the first place, each of the individual peaks in the frequency spectra could be tracked. With additional information, a particular frequency could be selected that represents the best-guess of the heart-rate. With the (incorrect) assumption that the heart-rate doesn't change particularly quickly, it would be possible to track the heart-rate in the presence of other noise sources that emerge and then decay away. Looking back to Figure 4.23, it could be envisioned how a tracking algorithm could maintain a 'lock' on the true heart-rate peak (the starred point), even when the artefact's amplitude increases above the PPG's. Such algorithms, however, must have additional processing involved to ensure that artefacts are not erroneously tracked.

Additional processing (post-processing) of the resulting heart-rate data (from the FFT method, and others) has not been researched further as many algorithms already exist, and are in use; it is instead the goal of this research to improve the PPG signal itself, such that algorithms that are used generate more accurate results. Ideally, to prevent algorithmic complexity from increasing, that will undoubtedly heighten the risk of something going wrong, the data that is fed to the heart-rate extraction tech-

niques should be as clean as possible to begin with. This will ultimately require development effort in the PPG extraction methodology, instead of relying of processing methods to correct for artefacts and errors.

4.5.3 Auto-Correlation

Auto-correlation is a technique that, like the frequency transformation methods, results in frequency information being obtained from the temporal input signal. This frequency information is, more specifically, related to the periodicity of the signal: when plotted, it is known as a correlogram. If a signal is highly periodic, then the autocorrelation will give a large response at this period, and multiples thereof, of the signal. By taking the location - or *lag* - of the first dominant peak of the correlogram, then the frequency can be obtained by simply taking the reciprocal of it (after scaling by the sample-rate).

Unlike the previously discussed frequency analysis, the resolution of the resulting period is not dependent on the amount of data used (the window size), it is instead the sampling rate that is important. With a sample-rate of 100 sps, the resolution of the period is 10 ms. As this technique produces a period (time domain) instead of a frequency (frequency domain), then the resolution of the frequency-domain heart-rate will vary with the heart-rate itself (as the reciprocal). Higher heart-rates will have a lower resolution, but even at 240 bpm (4 Hz), this is *relatively* small at just 10 bpm (see Table 4.1).

To obtain a more accurate value, two options are available; either increasing the sampling-rate, or re-sampling (and interpolating) the data. Both have the same effect, but re-sampling will only work up to a point - in a similar way to interpolating the FFT results - as there is no additional data gained.

Figure 4.27 demonstrates a naïve method of extracting the heart-rate.

Period	Frequency	Heartrate	Delta
240 ms	4.167 Hz	250.000 bpm	↑ 10.000 bpm
250 ms	4.000 Hz	**240.000 bpm**	↑ 9.231 bpm
260 ms	3.846 Hz	230.769 bpm	
990 ms	0.990 Hz	59.400 bpm	↑ 0.600 bpm
1000 ms	1.000 Hz	**60.000 bpm**	↑ 0.600 bpm
1010 ms	1.010 Hz	60.600 bpm	

Table 4.1: *Example heart-rates for given periods (periods are separated by 10 ms, a time achieved with a 100 sps sampling rate).*

In this example, the predominant period of the 5-second window of data occurs at 86 samples (0.86 seconds at 100 sps). This yields an 'average' heart-rate of 69.8 bpm over this window.

Of course, as this technique still requires multiple pulses to operate correctly (with an absolute minimum of two), the window size must still be taken into account. Additionally, less naïve methods of heart-rate extraction (as explained shortly) require much larger window sizes in order to exploit redundancy in the output. Ignoring this for a moment, the minimum window size (which will give the highest responsiveness), is the duration of two pulses at the lowest expected heart-rate. At 40 bpm (or 1.5 seconds per beat), the time for two full cardiac cycles is 3 seconds, in this case. For the examples given here, the window size has been set to 5 seconds, to match the previous (FFT) methods.

Using the first period 'peak' in the correlogram as the means of determining the heart-rate, however, has a major flaw. Short-lived artefacts such as those caused by sudden displacements of the sensor (it being knocked) will momentarily cause the first peak to be shifted to a new 'lag', or suppressed to the point that it is no longer a 'peak'. This happens because the short-duration artefact does not correlate well with the neighbouring pulses in the window. However, as other pulses will still exist, when the autocorrelation shifts one of the data-sets by *two* PPG periods, then the correlation is suddenly high again. This 'redundancy' in the output from

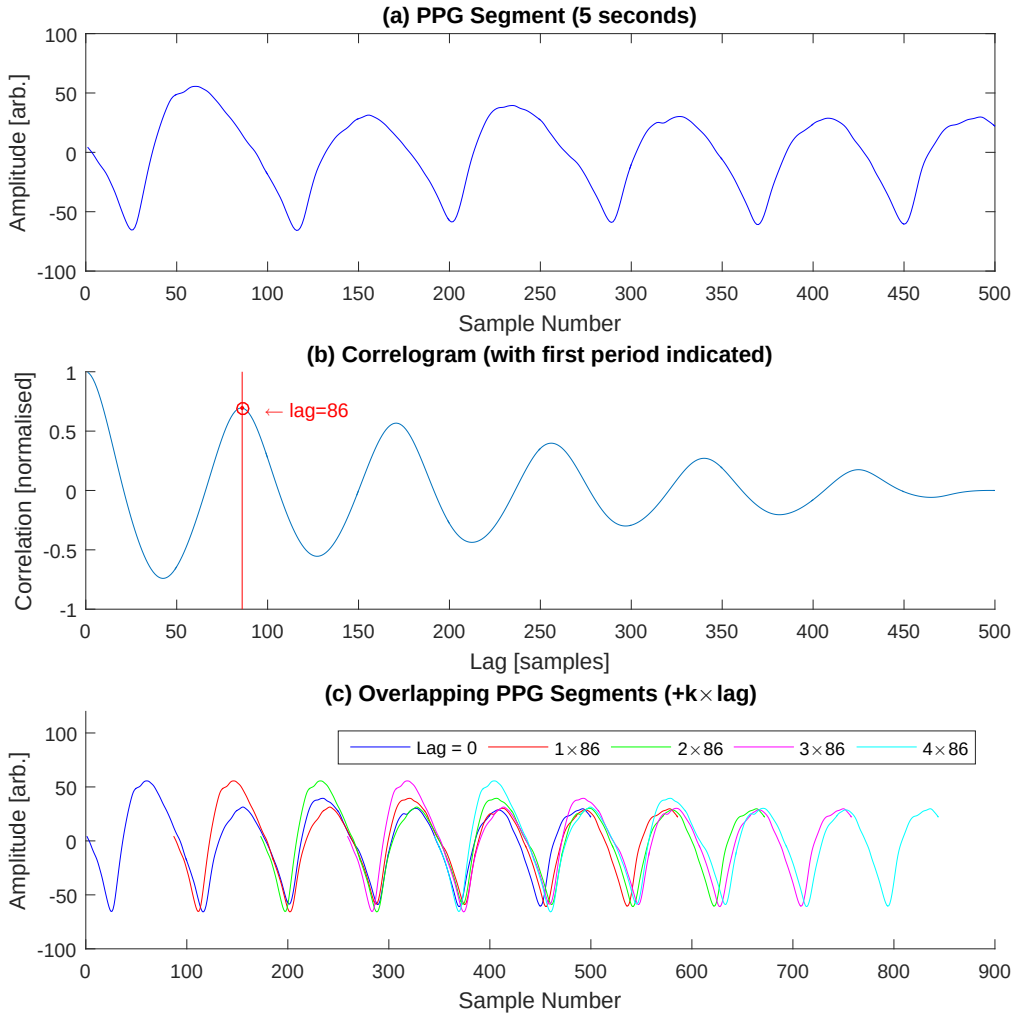

Figure 4.27: *Example illustration of a 5 second PPG window with its correlogram. The first peak on the middle graph (marked with a lag of 86) shows the offset that provides the largest correlation (this is represented as the red line on the lower graph). As the signal is highly periodic, it will align with itself well with all integer multiples of the lag.*

the autocorrelation can be used to predict the average period of the PPG without the existence of the first correlogram peak, and is only possible with a window that encompasses more than two complete pulses. The greater the number of pulses the window encompasses, the more peaks in the correlogram exist that all *effectively* represent the same information. The distance between these peaks is a good approximation to the position of the first peak.

Naturally, with larger window sizes, the output from the processing will have a larger delay. Also, the 'temporal resolution' of the output will be

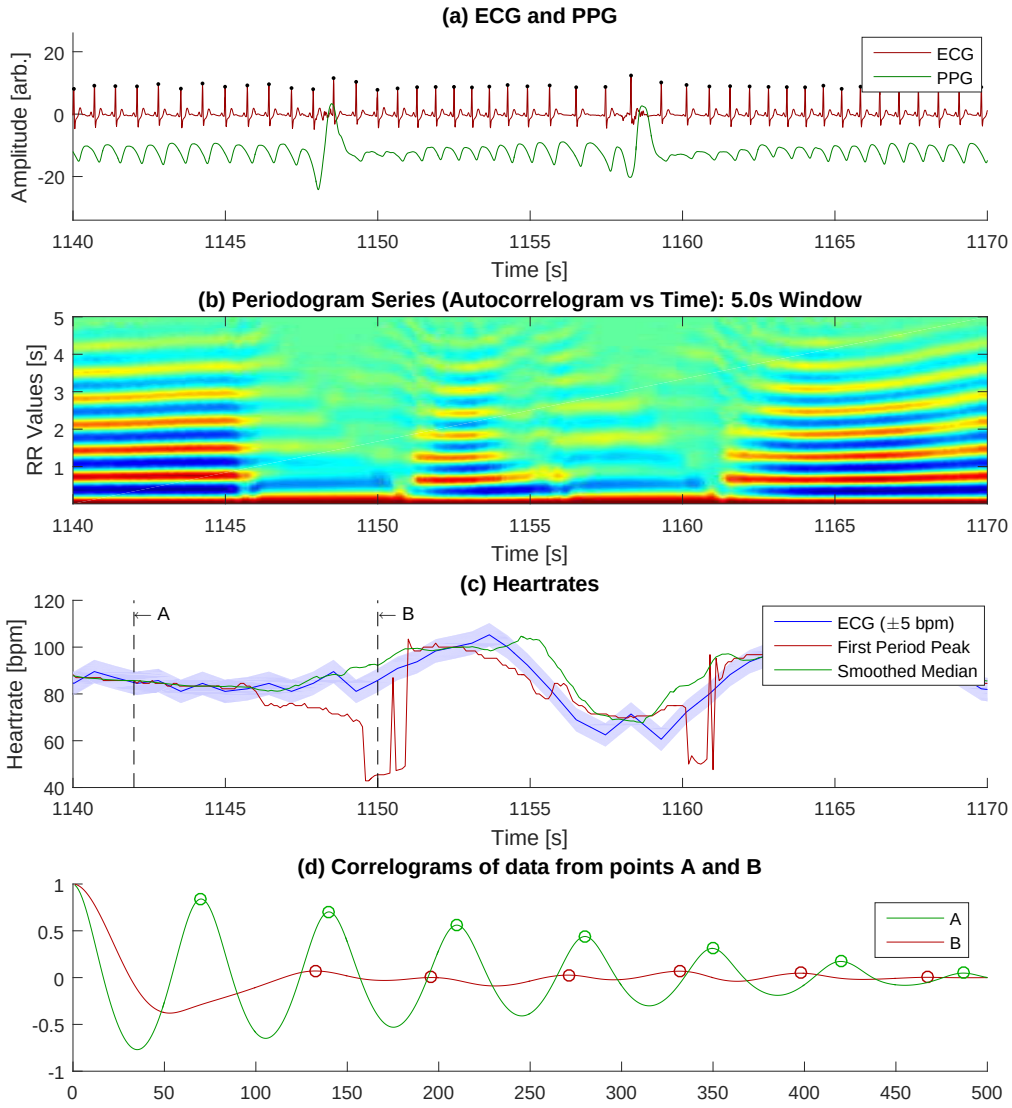

Figure 4.28: *Demonstration of the difference between using the first peak and the median of the differences. For the regions where the artefact-like disturbance exists, the periodicity is still clearly visible in the periodogram (graph (b), a collection of correlograms in time). Although weaker, the peaks are still visible in the lower graph (line B).*

degraded: the calculated heart-rate will not follow the true heart-rate as accurately, instead it will appear *smoothed*. Conversely, smaller window sizes will provide a much more accurate and timely result.

Figure 4.28 shows an example of two short-lived arrhythmia (bradycardia, which have similar characteristics to small artefacts) causing disruptions to the detected heart-rate when just the first 'peak' in the correlogram is used (i.e: the naïve method). Line 'B' in the plot (d) shows the correlogram of a window at one of the arrhythmia (bradycardia) missing the first

peak (the peak is still there, but has no negative-going inflection, making it invisible to the standard peak-finding algorithm used). However, as the PPG pulses still exist in the surrounding signal, other peaks exist in the correlogram that - as discussed previously - allude to the heart-rate.

By taking the median of the differences between each of the adjacent peaks' lags, a good approximation of the PPG period can be obtained. The median was chosen instead of the mean as - depending on the magnitude of any present artefacts - other peaks can be effected: either by being suppressed or by having additional peaks inserted between them; and it is the 'middle-most' difference that is attributed to the PPG period.

4.5.3.1 Results

Two examples of the heart-rates detected from PPGs are presented next. Figure 4.29 shows a region of clean, *well behaved*, PPG with one small artefact that was likely caused by the sensor being 'knocked'. Auto-correlation accurately (within the ±5 bpm indicated by the light blue band) follows the actual instantaneous heart-rate as determined by the ECG. The small artefact, despite being more than three times the amplitude of the surrounding PPG, has left the output mostly unaffected.

One notable feature of the resulting heart-rate plot is that it is much smoother than that of the ECG's instantaneous heart-rate. This is due to the aforementioned windowing effect whereby each heart-rate is calculated from the surrounding data (a 5 second window was used in this instant). This also has a side-effect of biassing the heart-rate to the local mean; between 1280 seconds and 1290 seconds in Figure 4.29 the heart-rate from the auto-correlation method appears to move towards the mean (as if it were low-pass filtered). If the heart-rate were to suddenly - but momentarily - change, then this technique may under/over-estimate the heart-rate for that period. A smaller window would improve the response time, but

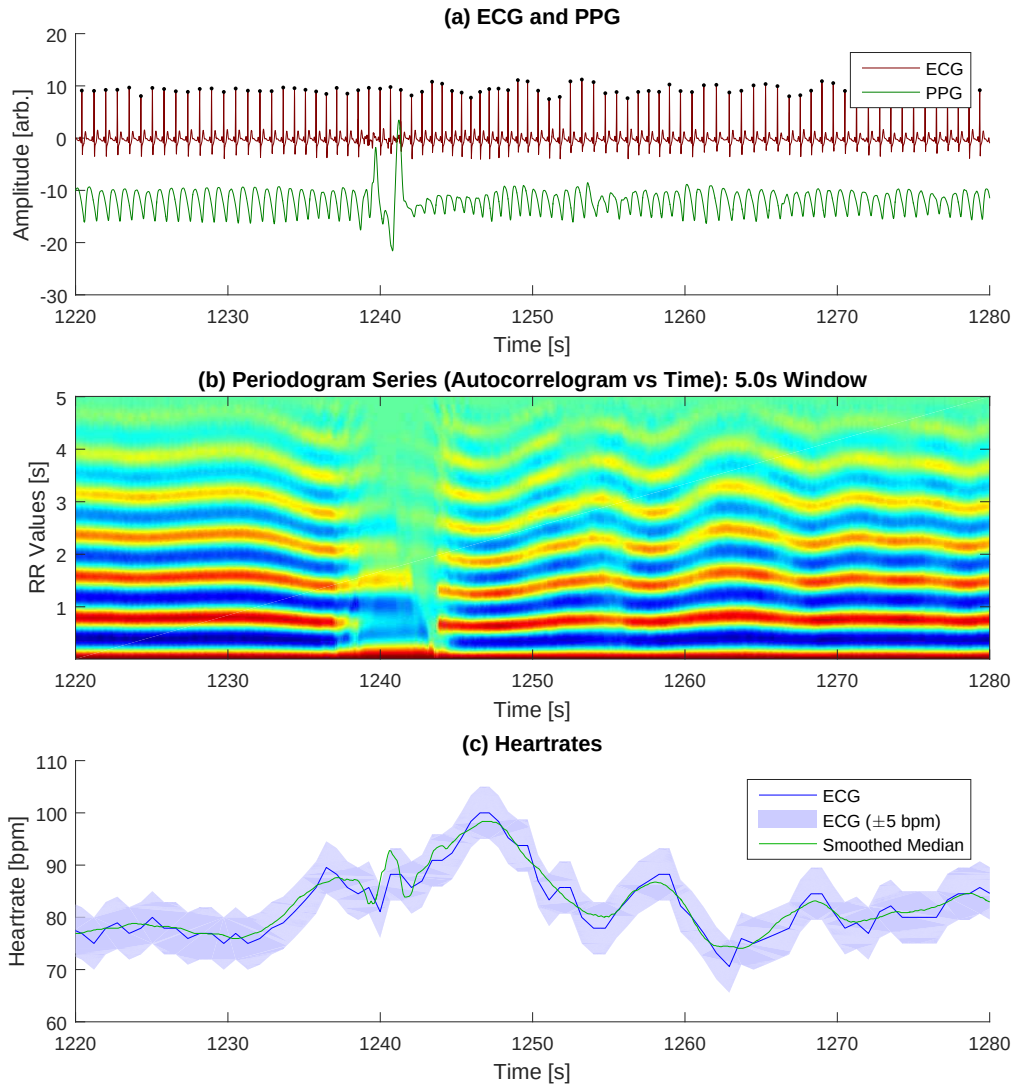

Figure 4.29: *A clean section of PPG (bar one small artefact), showing the effectiveness of the auto-correlation method.*

would leave the technique more susceptible to artefacts.

Figure 4.30 shows a motion artefact corrupting the signal. This artefact, unlike the previous one, is highly periodic (albeit short-lived). This periodicity is visible in the periodogram (the central plot) as a lower frequency ripple, and shows the downside to the technique. For repetitive (periodic) artefacts, auto-correlation will result in them being detected instead of the PPG. An algorithm similar to that used in the previous FFT method could be used to correct for such anomalies, but where information is missing, the best the algorithm could do is 'fill in the blanks' with an estimation of the heart-rate (a dangerous prospect in clinical settings).

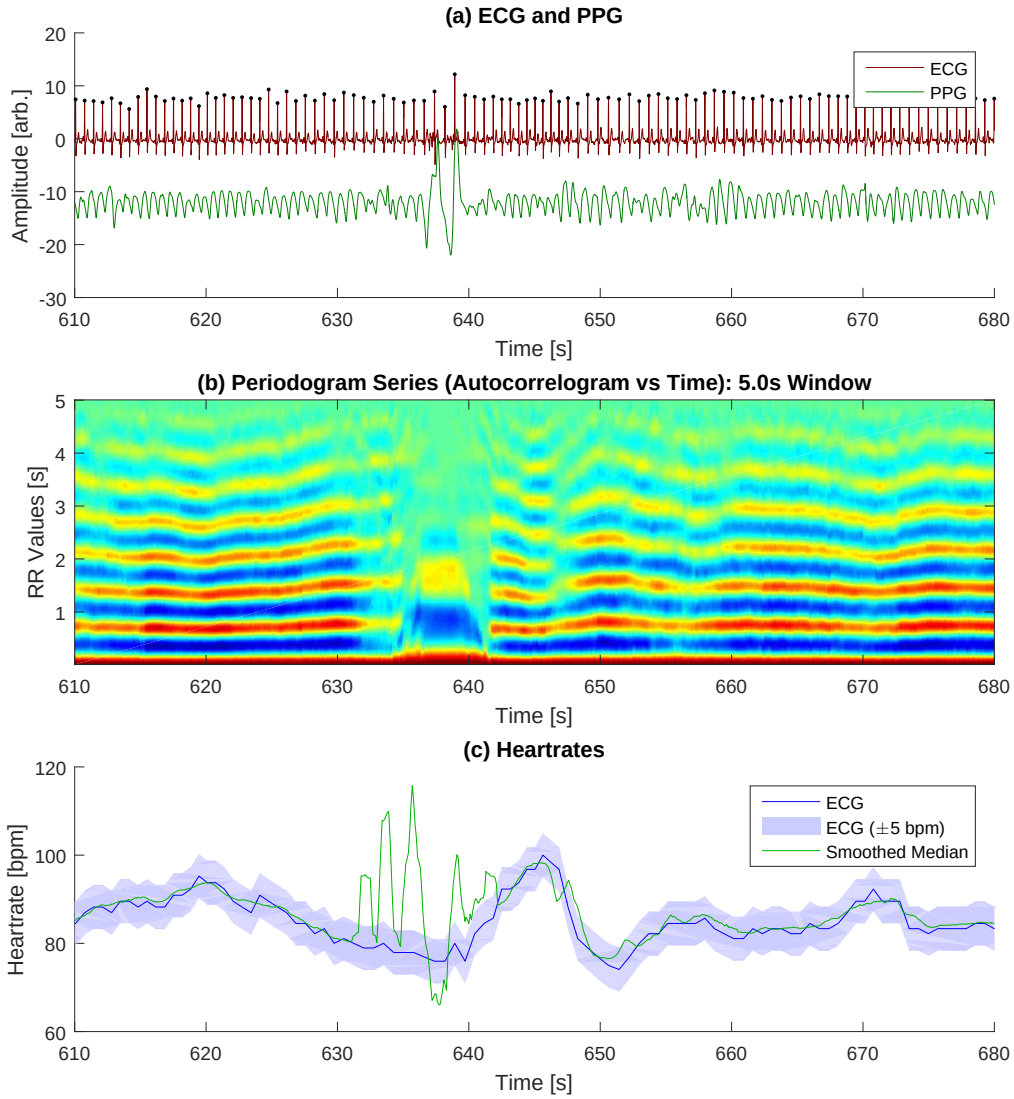

Figure 4.30: *An example of a larger artefact (the duration of which is similar to the window length, and is periodic). In this case, the detected heart-rate intermittently jumps to other values (the period of the artefact).*

4.5.3.2 Summary of Autocorrelation Technique

Unlike the previous method discussed (the frequency analysis technique) which required an additional step to correct for errors in the processing, the autocorrelation method presented here does not have a similar processing step. This is due to the higher resilience of this method to artefacts in the signal, and a more reasonable (less pronounced) error on the output when artefacts are significant. This is most likely due to the effect of using the time domain data instead of transforming it into the frequency domain,

where the power/amplitude of the data has more of an effect. The auto-correlation technique therefore allows for greater reliability when detecting the heart-rate compared to the FFT method.

A quantitative comparison of the effectiveness of this technique compared to the 'peak-detection' algorithm is not easy to make as there are multiple factors that involve knowledge of the application they will be applied to; for example, when newborn resuscitation is considered, the peak-detection algorithm will provide the faster response (the more instantaneous heart-rate), but will be more susceptible to errors without additional processing (meaning extra time). The auto-correlation technique, on the other hand, will have an inherent delay before the values are output, but will provide a greater resilience to artefacts - at the expense of reduced temporal resolution.

4.6 Indirect Quality Measurement

Whilst a detected heart-rate that is 'out-of-range' from what is expected (either by it being infeasibly low[12], or impossibly high), is suitable for rejecting PPG data, there will be times when such heart-rate extraction algorithms will output a reasonable value when in fact no clean PPG exists. The following techniques present methods that achieve rejection of a different quadrant of the truth/false positives/negatives grid. These techniques are suitable for rejecting data that does not conform to what would be considered a 'normal' PPG, at the risk of rejecting good 'quality' signals.

[12]For the given patient/participant, assuming they are healthy. If their heart-rate really did drop too low, then the 'out of range heart-rate' would be reason for a doctor to perform a manual check, and intervention, if needed.

4.6.1 Out-of-Band Analysis

For previous methods, the raw PPG signal from the recording devices have been filtered to attenuate frequencies that lie outside of the typical heart-rate range. As previously discussed, by doing this, information about possible artefacts may be lost (as any filtering will remove at least some information). This section looks at the analysis of 'information' that is out-of-band of the typical PPG.

If, for example, a participant's PPG was recorded in ideal conditions (no motion, constant light levels, etc), then the PPG that is extracted would primarily contain information about the pulsatile blood flow (other components related to breathing and vascular constriction - Meyer waves - would also exist), and would not contain artefacts. The frequency components associated with the PPG are well-defined and are bounded - as previously discussed - by frequencies at the minimum heart-rate and at approximately three times the maximum heart-rate (giving two harmonics). More harmonics exist, but are of a sufficiently low amplitude to be ignored.

Signals that directly interfere (corrupt) the PPG are those with frequency components that lie within this band. However, such signals are unlikely to lie ***entirely*** within this band; there will be components that are outside of this range (out-of-band).

An observation made of typical raw signals from PPG devices is that artefacts with frequencies *far* above the PPG's natural range (25 Hz or greater) are likely due to artificial lights 'flickering' - 60 Hz from PC monitors, 100 Hz from mains lighting, etc - and can be excluded from this analysis as these will not have components in the PPG range[13].

It is therefore reasonable to assume that any frequency components that are close to the expected heart-rate range (but are outside of it) are most

[13]This is assuming that there is no aliasing between the sampling frequency and the noise/artefact source.

likely due to artefacts. Analysis of these bands can yield a value that can be an approximate representation of the signal 'quality'.

This technique can be achieved by one of a multiple of methods, but the simplest to understand (which is also the least computationally expensive) is by using different filters of varying band-widths and centres. To obtain a filter that extracts the signal with frequency components either side of the heart-rate range, two band-pass filters could be implemented and combined to create a single filter with the necessary characteristics. Whilst perhaps being the most intuitive, a more convenient method uses the overlap between two filters - one of which already exists in the standard set-up.

The two filters that have been used are illustrated in Figure 4.31. The 'PPG Band' band-pass filter has been used previously in all heart-rate extraction techniques to reject out-of-band frequencies.

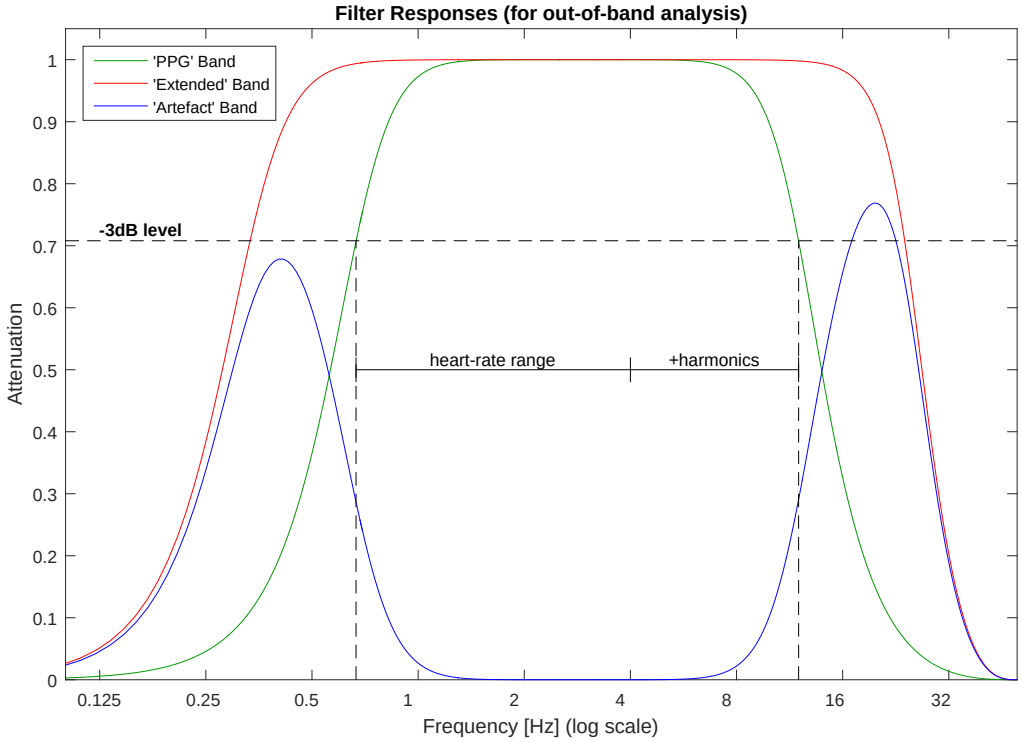

Figure 4.31: *The 'extended' band (in red) has frequency cut-offs that are half the minimum and twice the maximum of the PPG's. Note that the frequency axis is following a log-scale to improve clarity.*

The 'artefact' band(s) that are visible are created simply by taking the difference between the two filter responses. In the actual processing, the same signal (the PPG) will be filtered independently by each of the two filters and the results will then be subtracted - the 'PPG' range from the 'extended' range. This has the same effect as a specialised filter as previously described. Another advantage of taking the difference between filtered signals is that no information is lost; if two different filters were used, for example, to extract frequencies below and above a given cut-off frequency, then the sum of the two filtered outputs would *not* equal the original signal due to the filters responses not summing to one[14].

The difference between the original signal and the band-pass filtered signal can be considered as the 'residual' and is presented in Figures 4.32, 4.33 and 4.34 (a). The lower plots of the same figures show the actual filtered signals.

[14]Although it is possible to design filters to do this, simply setting the cut-offs of a low and high-pass filter to the same value will not yield the correct result.

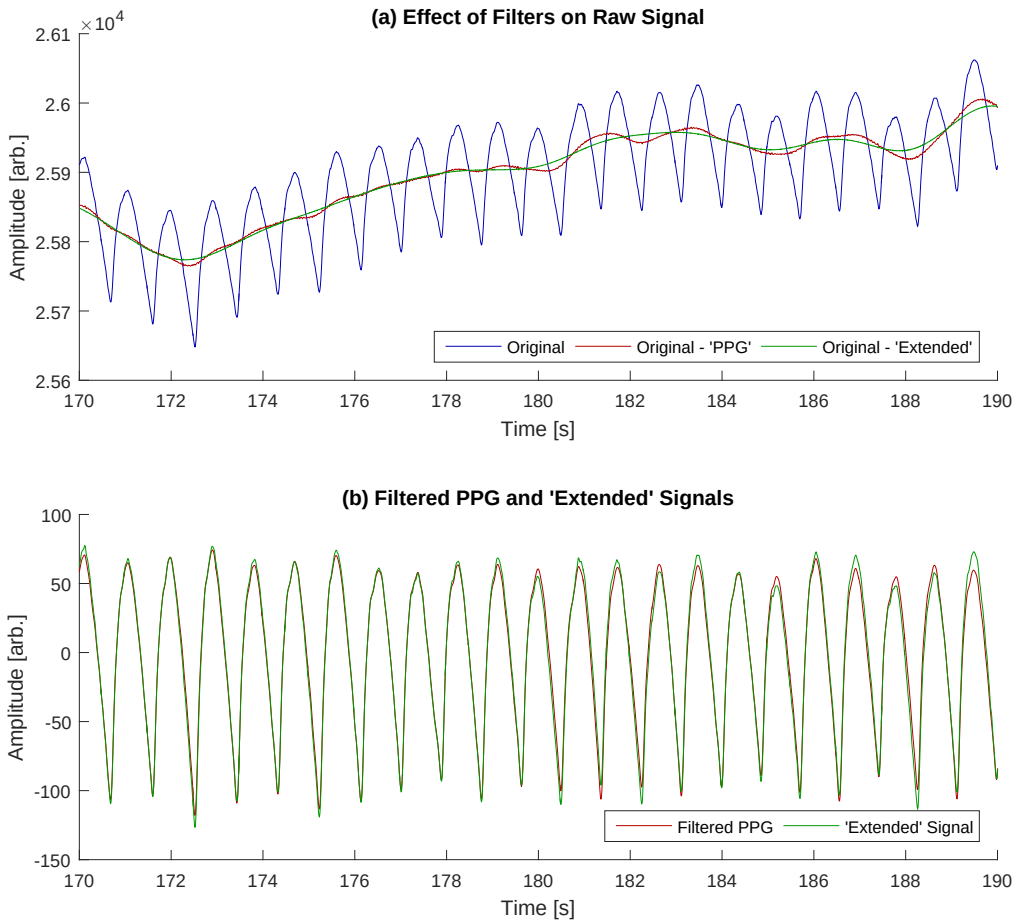

Figure 4.32: *A 'good' clean signal. The difference between the two filtered signals is minimal, illustrating that there is little to no signal that is present in the 'artefact' region.*

To quantify how 'corrupt' the signal is, the amplitude of the filtered PPG can be compared with that of the artefact's. This yields what is effectively a 'PPG-to-artefact ratio'. In this particular case however, as it provides a more clear result for illustration, the reciprocal is used (giving the artefact-to-PPG ratio).

As all signals that are used at this point have been band-pass filtered (i.e: including removing low frequency components and 'DC'), then the signals themselves sit around zero. A simple envelope detector (algorithmically based off the circuit in Figure 2.6) is therefore sufficient to compute the envelope, and therefore a rough estimate of the signal's amplitude. In all the following cases, the envelope of the absolute value of the signal is used (the envelope of the PPG signal has a tendency to extend above the

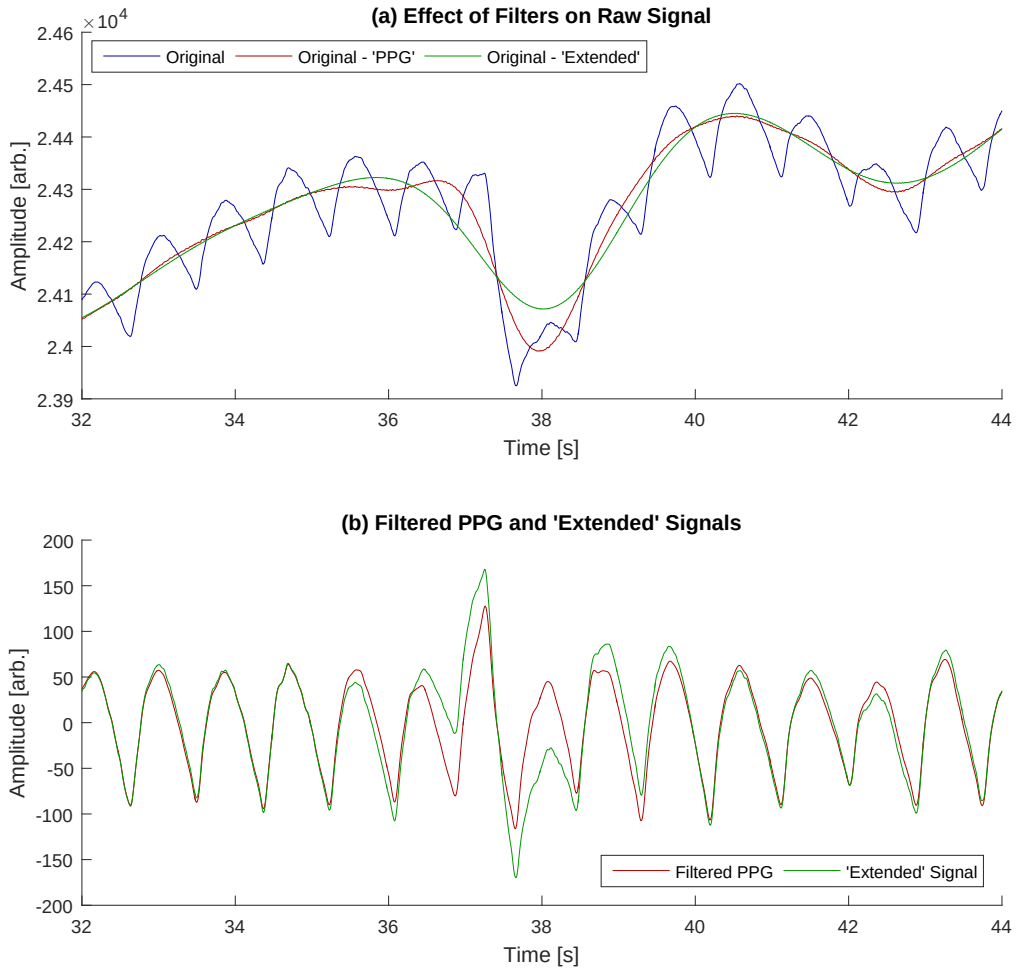

Figure 4.33: *A small artefact causing a sudden change in the signal's base-line. Whilst the PPG filter almost entirely removes this, the extended bandwidth signal doesn't - this causes the difference that is used later.*

peaks of the waveform due to the troughs being more prominent than the peaks despite locally being zero-mean).

In Figure 4.35, the small spike in the ratio plot at 205 seconds is due to a sudden decrease in the amplitude of the PPG with a small increase in the artefact's amplitude. For the sample set of data that was used, when this method yields a ratio of less than 2.0, the PPG is still easily discernible (a qualitative assessment based of the observed results).

Figure 4.36 illustrates a signal with a larger artefact (one with approximately twice the amplitude). The ratio associated with this event is still small, but is larger than the previous case, showing the sensitivity of the technique. The amplitude of the PPG is approximately 100 'units' at the

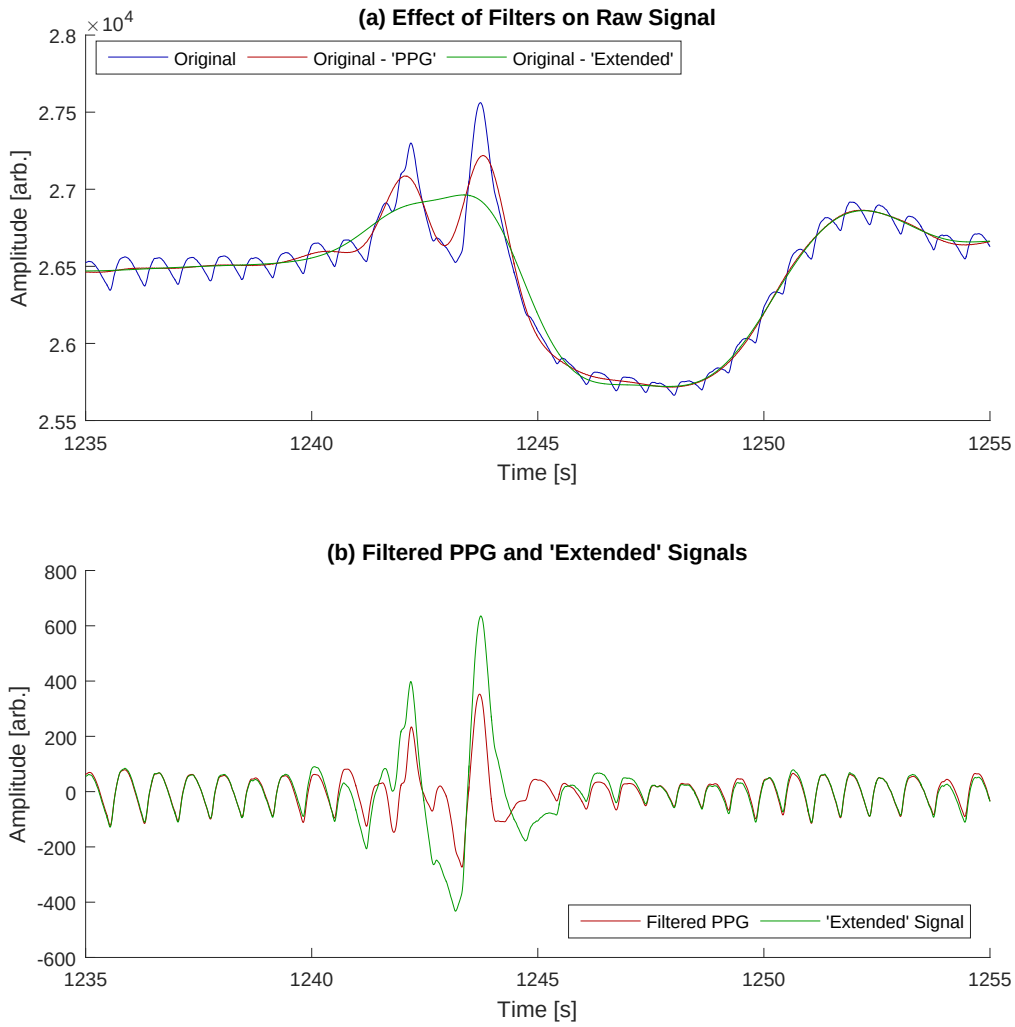

Figure 4.34: *The artefact in this signal completely corrupts the PPG. The difference between the filtered PPG and extended bandwidth signals is large.*

event, not much more than the surroundings, whilst the amplitude of the artefact signal increases by more than five times. Finally, in Figure 4.37, the ratio of 'artefact amplitude to PPG amplitude' has increased during the artefact event to more than eight.

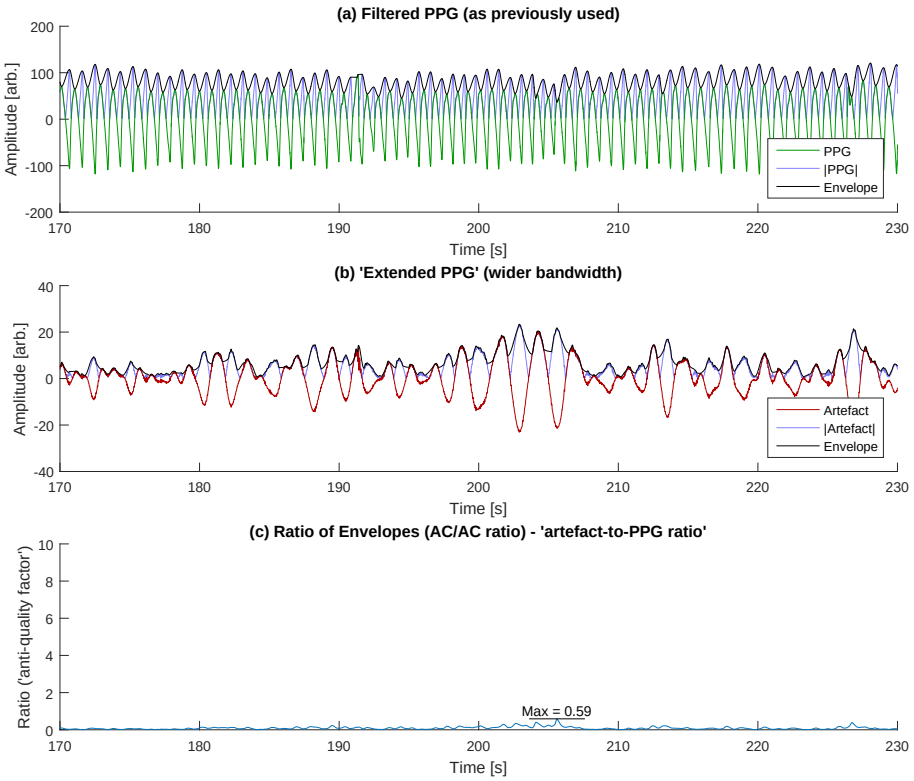

Figure 4.35: *A section of PPG with no significant artefacts present. The artefact-to-PPG ratio lies entirely below one.*

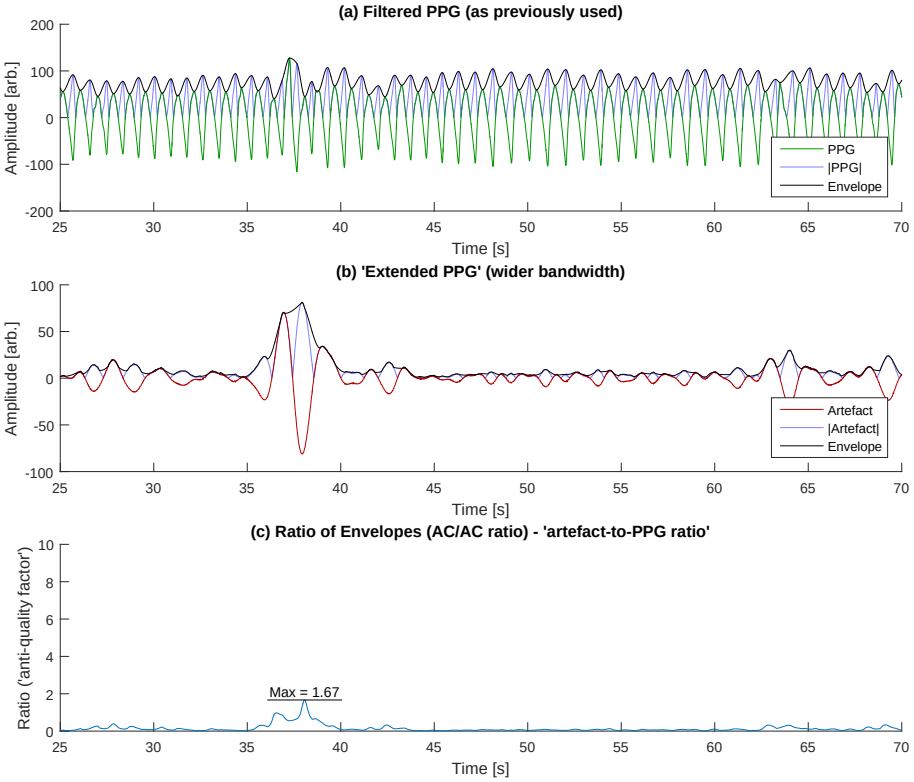

Figure 4.36: *A section of PPG with a small artefact.*

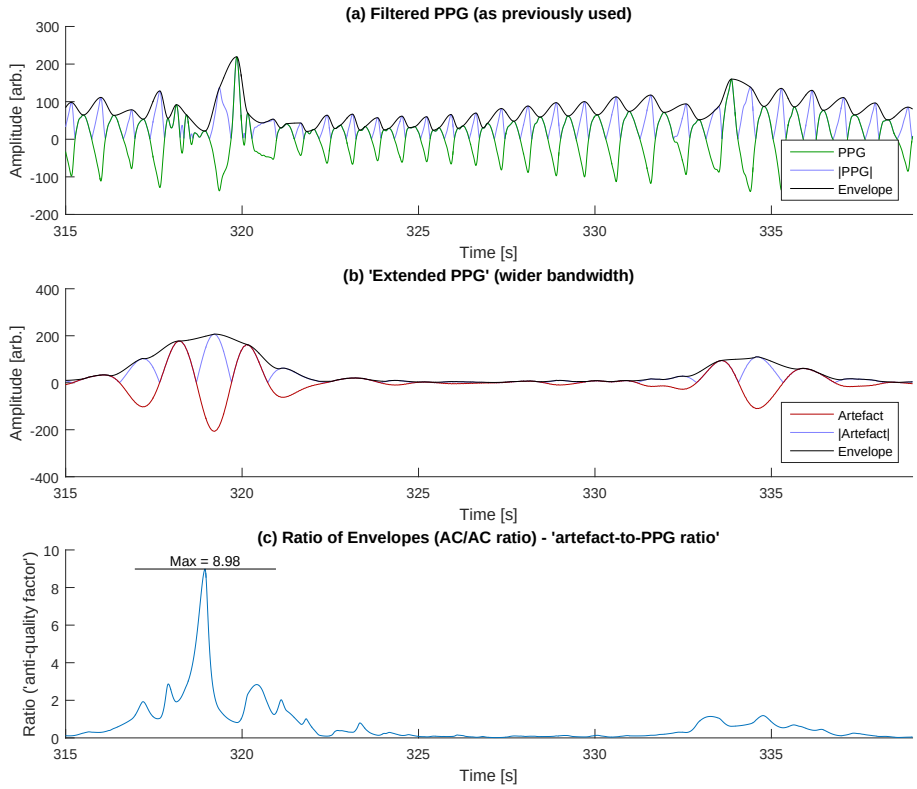

Figure 4.37: *The artefact present in this signal has corrupted the PPG completely for a short duration. The artefact-to-PPG ratio has now increased significantly. The artefact band shows activity 'ahead of time' simply due to the application of the filter (it was run forward and in reverse to remove group delays).*

4.6.1.1 Summary of Out-Of-Band Analysis

As discussed previously, neither this particular algorithm (nor the ones to follow) are useful in their own right as there is no heart-rate output. Instead, they should inform other algorithms about the quality of the signal such that they can proceed accordingly - there is no point in attempting to calculate a heart-rate, for example, if the signal 'anti-quality' (artefact-to-PPG ratio) as given by this method is more than, say, three; errors are otherwise likely to occur.

The primary limitation of this technique is the fact that if an artefact's frequency components were to lie entirely within that of the PPG's range, then the resulting ratio would not indicate the artefact's presence. This technique relies on any artefact that would potentially corrupt the PPG

having sufficiently spread out frequency components as to lie partially outside the heart-rate (plus harmonics) range. Whilst such an artefact would have to be highly periodic and sinusoid-like, it **is** possible for one to exist.

One possible way of improving the selectivity of this algorithm would be to adjust the weighting of the artefact bands (of the filter). Through observation, it is clear that the majority of artefacts encountered have more dominant lower frequency components [69]; the filters could be modified, therefore, to enhance the higher frequencies in the 'artefact band' to allow for higher sensitivity to those artefacts.

4.6.2 Envelope Analysis

When the raw PPG signal is band-pass filtered using a 'heart-rate-range' filter, frequencies that are out-of-band are removed as previously discussed. One of the down-sides to this approach is that some artefacts may not extend outside of the natural heart-rate range at all. For these artefacts, additional analysis of the in-band signal must be performed. However, as looking at the frequency components could quite easily lead to confusion over which components belong to the PPG and which to the artefact (this is assuming one or both actually exist in the signal), then the baseline and amplitude of the signal - or more specifically in this case, the envelope - can yield the desired results.

The two envelopes (upper and lower) of the signal in this case represents the sum and difference of the natural baseline and the amplitude modulation of the signal, where the PPG is the 'carrier'. This signal has frequency components of its own, of course, but will be close to the PPG's fundamental frequency due to the mixing effect. For artefact-free PPGs, the envelopes of the signal will naturally move apart and together in unison (modulation), as well as 'up and down' (perturbation) as a result of breathing and vascular constriction (Meyer waves etc). The amount that

they deviate from each other, and from their natural baseline, however, will often be minimal and tend to have much longer periods.

Artefacts too can cause a similar response, but due to their mechanical nature, they will cause compression and stretching around the measurement site, leading to exaggerated deviations in both the envelopes' modulation and perturbations. These larger 'artefacts' appear as larger modulations and perturbations in the signal. In some cases, a large baseline shift can be caused by factors that do not affect the PPG signal itself; similarly, the PPG's peak-to-peak amplitude can deviate suddenly without an artefact existing. When deviations occur at the same time, then the chance of motion being the cause is higher.

By taking the differentials of the envelopes' sum and difference, an idea of how 'turbulent' they are can be obtained. By taking the product of these two values, a good estimation of where artefacts are likely to be will be formed. The *product* was chosen in this case as the resulting value would only grow if both of the differentials were sufficiently large. If one was small, for example, then the other would have to be considerably larger for there to be an equivalent 'artefact factor'.

Figure 4.38 demonstrates the effectiveness of taking the product of the two differentials, as opposed to using each individually. Most of the time, at least one is close to zero; only when an artefact occurs do they both increase in magnitude. The short and brief fluctuations that occur in the $\Delta amplitude$ signal (visible in the middle plot) are due to the small dicrotic notch in the PPG signal causing a small step in the envelope[15]. Also visible in the figure is the natural variation on the amplitude and perturbations of the PPG (top graph). Even though the baseline changes quite dramatically, the differential is quite small due to the long periods.

[15]The developed envelope 'acquiring' algorithm, presented in Appendix I, uses a lossy peak detector. The shallow gradients at the top of the PPG pulse allow for the envelope to 'drop' into the space above the dicrotic notch.

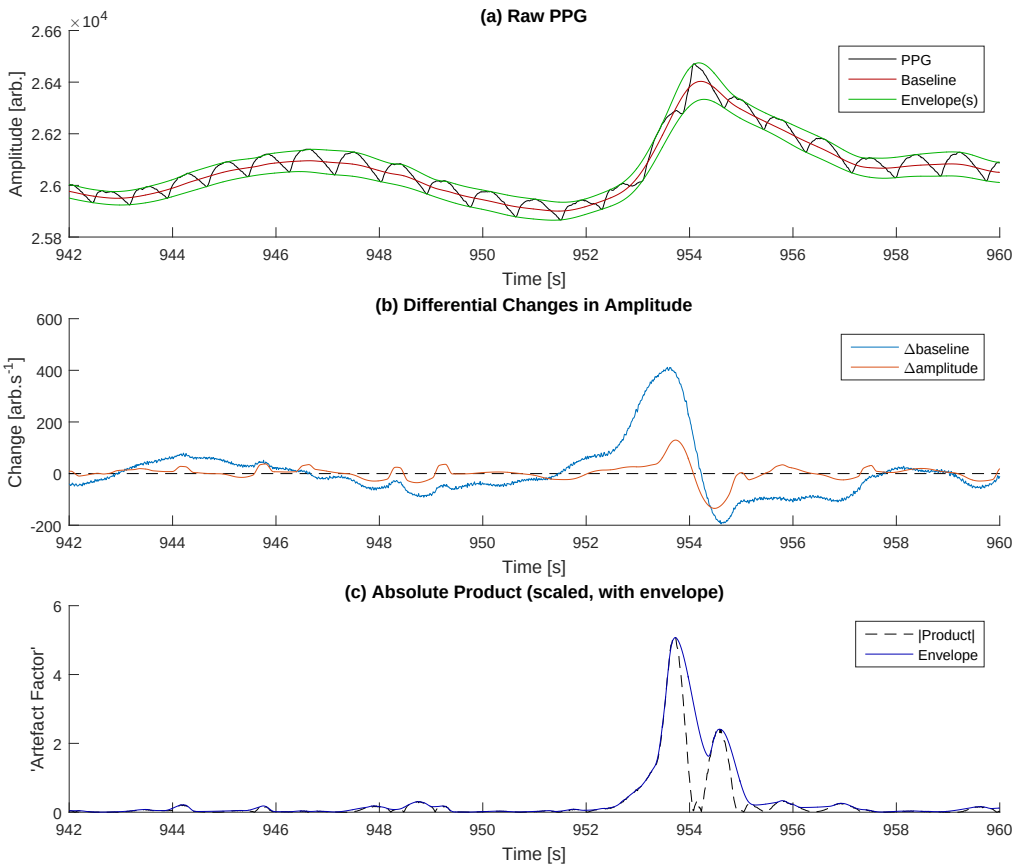

Figure 4.38: *Example of a small artefact in an otherwise 'clean' signal. The individual differentials (the baseline and the amplitude) can swing around zero by a reasonable amount, but only when both have a larger magnitude does the product (the 'artefact factor') increase substantially.*

Despite the middle graphs in Figures 4.39 and 4.40 having a vertical scale of 0 to 10, the actual 'artefact factor' is arbitrary and has no meaningful units. Just like in the 'out-of-bounds' method discussed previously, the larger the value, the greater the chance of the corresponding region in the raw signal having a corrupting artefact. It is therefore important to realise that - as a result of this - the value will vary from one recording method to another, and thresholds must be adjusted accordingly. For the system in which the recordings were made, an 'artefact factor' of more than two is observed to be a reasonable threshold.

Figure 4.40 shows a PPG signal that has been deformed by an artefact. At around 368 seconds, the PPG in the top and bottom graphs show a sudden change in amplitude at the same time resulting in the expected

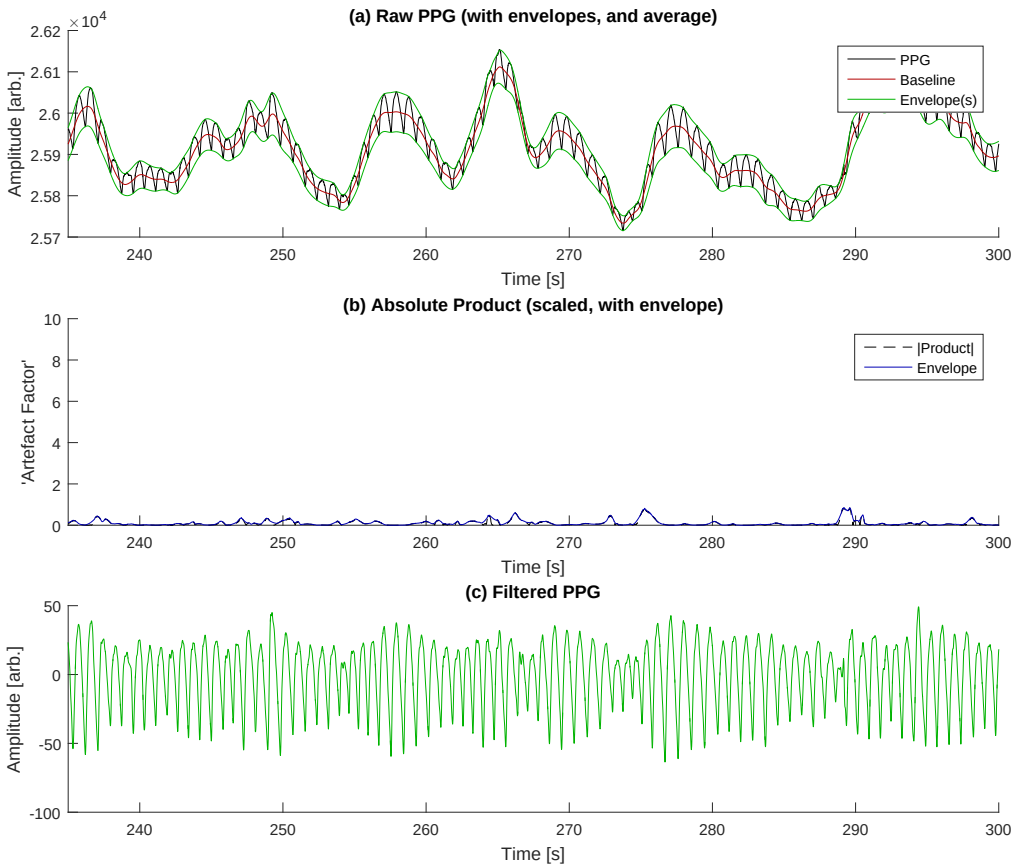

Figure 4.39: *A good 'clean' section of PPG. The baseline in the upper graph clearly swings quite wildly: by more than five times the PPG's amplitude; and the amplitude itself varies by more than 50% as is visible in the lower graph. Despite this activity, the 'artefact factor' (middle graph) remains consistently low.*

increase in the 'artefact factor'. As previously discussed, the noticeable change in amplitude at 358 seconds does not cause such a response as it is both not as significant as that of the artefact, and it is not coupled with a large change in the baseline.

Although this technique is relatively simple, it has been shown to detect the locations of artefacts with a relatively high effectiveness. Again, on its own, it may not be completely reliable, however, when coupled with other algorithms it may produce superior results.

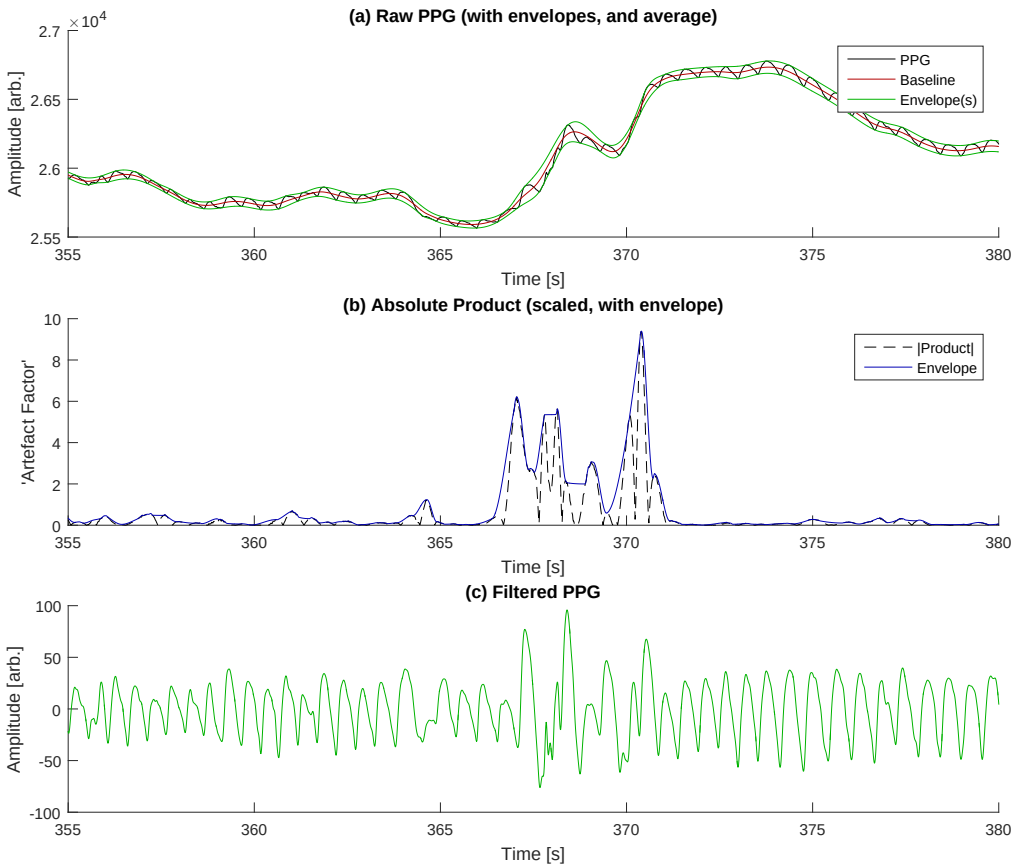

Figure 4.40: *A partially corrupted signal showing both the PPG amplitude and baseline varying together (in the top graph).*

4.6.3 Normalisation Effort

To further explore the possibilities of using the signal's envelope as an indication of the location (and severity) of artefacts, the amplitude (peak-to-peak) of the *filtered* PPG signal can also be processed. The filtered PPG, using the previously defined PPG filter, will of course remove the baseline leaving the frequency components that are associated with the heart-rate (including harmonics).

A simple 'decay' style envelope detector, which behaves much in the same was as the previously presented AM demodulator, can be used to extract an envelope that is insensitive to the noise in the signal. This envelope detector is run both forwards and in reverse, then the maximum of the results is taken; this is to make the envelope symmetrical around large peaks. The 'decay constant' that is used is dependent on the sample

rate, the expected heart-rate range and is given by the formula:

$$decay = \lambda^{(\frac{bpm}{60 \times sps})} \tag{4.1}$$

where *sps* is the sample rate of the data, *bpm* is the minimum heart-rate that is expected, and λ is a constant indicating the decay 'half-life value': a value of 0.5, for example, means that in one period of the lowest heart-rate, the decay value will have decayed to half of what it started at. For a more aggressive envelope decay, the constant can be set lower (see Figure 4.41).

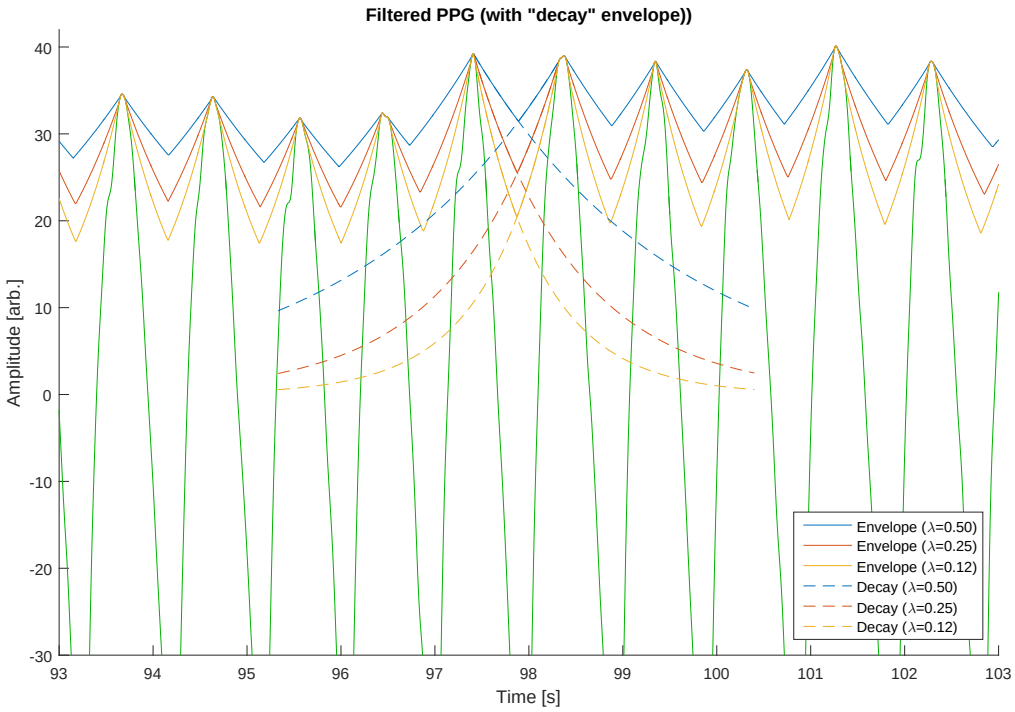

Figure 4.41: *Illustration of three decay 'half-life values'. The lower the 'λ', the faster the decay. All values of λ will result in the 'envelope' decaying to zero, but the rate will be different.*

This method of obtaining the envelope yields a 'signal' that is inherently 'spiky' and does not directly follow the PPG's peaks. When normalising the signal using this envelope (as can be seen in Figures 4.42, 4.43 and 4.44 - central two plots), the shape of the PPG is not preserved. The resulting signal is normalised between ± 1; by normalising it between these

two values instead of between zero and one, the signal remains symmetrical around zero, making processing easier (particularly the second pass of normalisation that follows).

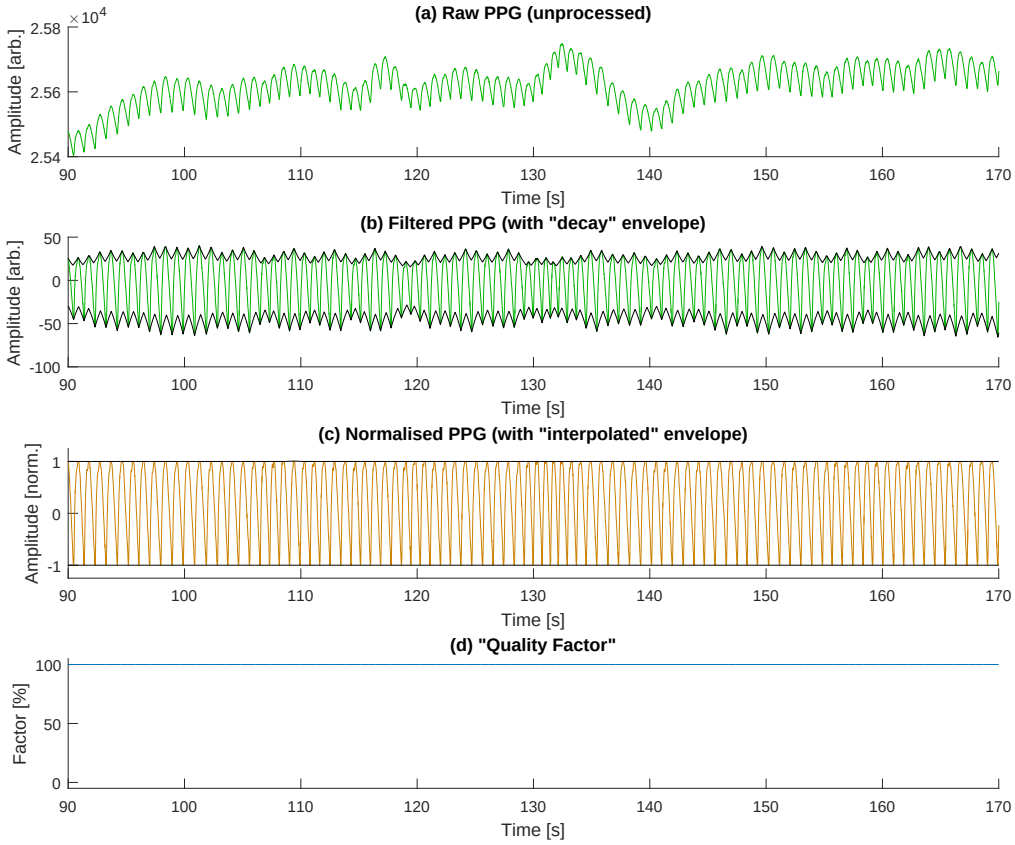

Figure 4.42: *A clean PPG signal with little variation. The resulting normalised signal has been normalised 'fully', resulting in no 'error'.*

Figure 4.42 shows what would happen if the signal contains no artefacts. The normalisation succeeds entirely resulting in the PPG signal lying between two 'tram tracks' at ± 1, with the peaks and troughs at these levels.

When an artefact exists in the signal, the envelope decay may not be sufficient to keep up with the rate of change of the peak-to-peak amplitude of the PPG. When this occurs, the normalisation process will 'fail', leaving the section of PPG containing the artefact not normalised. The '*effort*' that was put into normalising the signal did not pay off. To be more specific, when the artefact causes a sudden *decrease* in the PPG amplitude, the normalisation procedure will fail to increase the amplitude of the signal

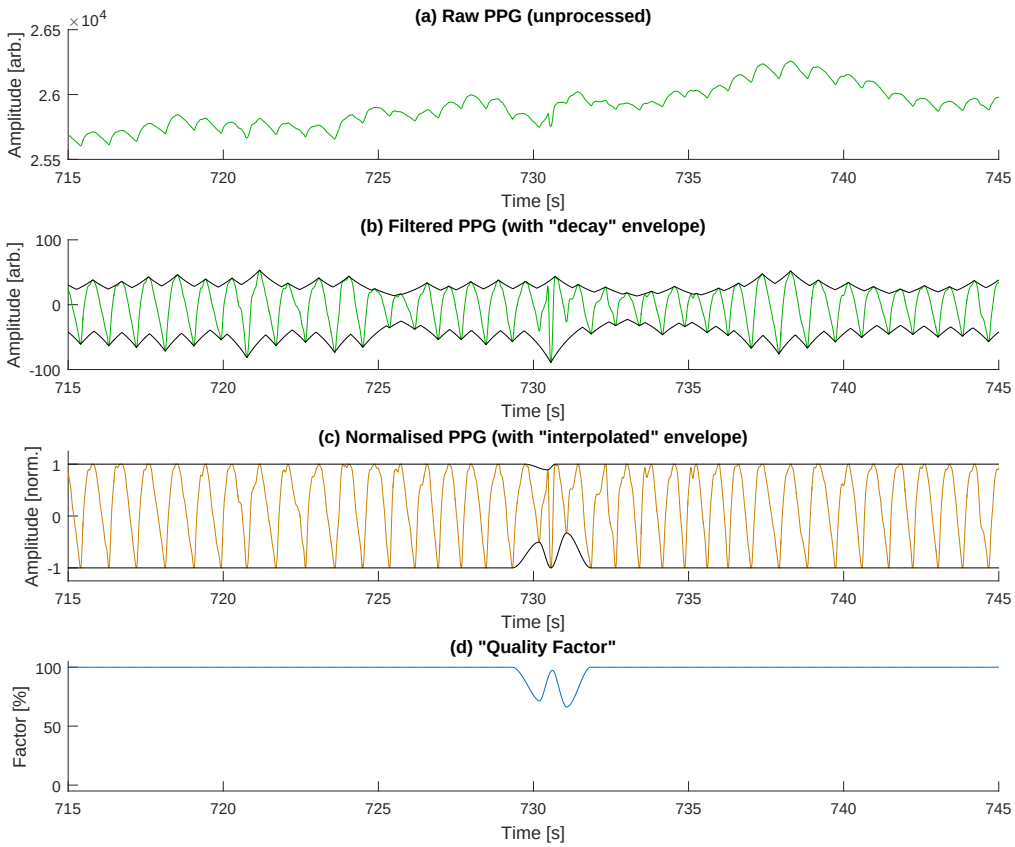

Figure 4.43: *A short-lived artefact with increasing amplitude causes to 'error' regions either side of it as the decay-style envelope detector cannot 'keep up' with the changing amplitude.*

back to the norm; however, when an artefact causes a signal's amplitude to *increase*, then the surrounding signal will not be normalised. Figure 4.43 shows what happens when the signal increases in amplitude suddenly, and demonstrates why the envelope detector is run both forwards and backwards (to make the envelope symmetrical about the peaks). Without the symmetry, only the 'falling edge' (the second error) would show.

The final step in this technique is to find the 'envelope' of the (partially) normalised signal in order to see where the previous envelope/normalisation step did not succeed. However, this time, instead of using a decay-style envelope detector, a simple thresholded peak detector (with interpolation) is sufficient. The peak detector is looking for the peaks and troughs of the normalised signal, ignoring any small fluctuations that may be caused by the presence of the dicrotic notch. MATLAB's `findpeaks` is sufficient for

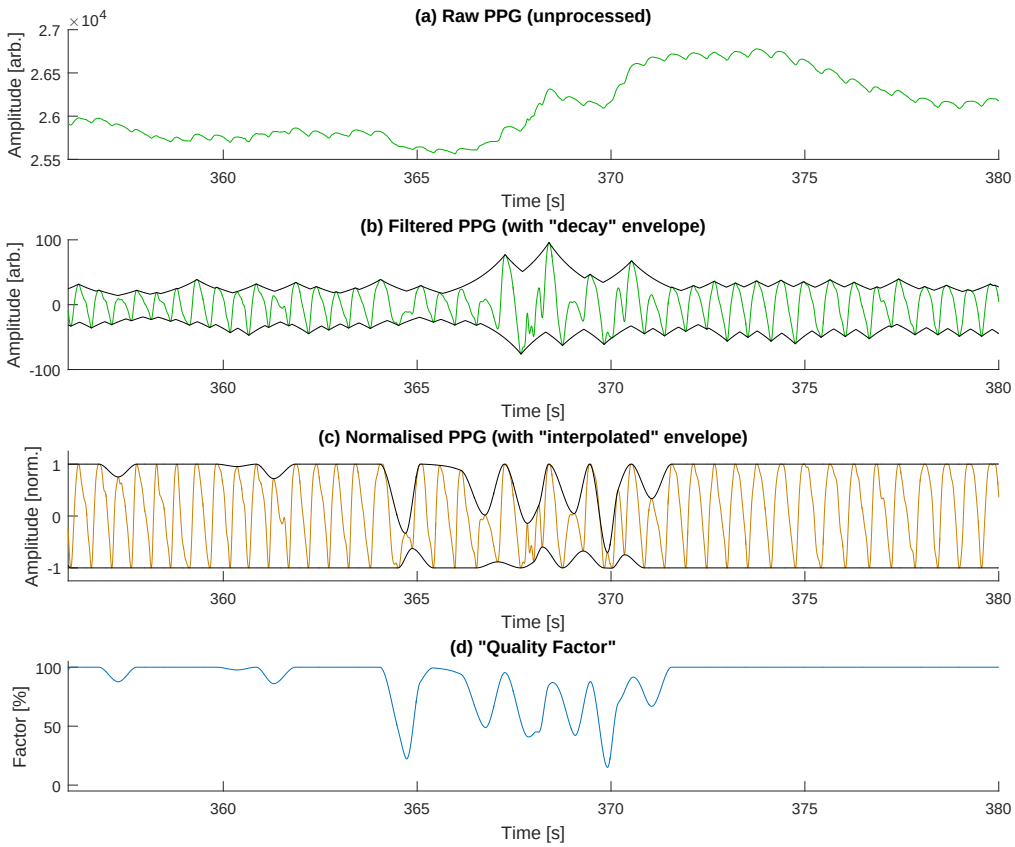

Figure 4.44: *A larger duration artefact created several low quality 'troughs'. A post-processing method could easily be applied to this signal to detect the envelope of the quality factor.*

this task with a 'minimum prominence' set to 0.25 (this is one eighth of the signal's amplitude as it was normalised to ±1). The resulting peaks and troughs can be interpolated to get the secondary envelope.

The difference between the upper and lower envelopes can then be scaled (to between 0% and 100%) to act as a 'quality' signal. A value of 100% occurs when the peaks and troughs are aligned exactly on the ±1 limits. A smaller value indicated some irregularity in the PPG's amplitude. A value of 0% is unlikely to occur due to the peak detection's 'prominence' rejection removing peaks that are smaller than an eighth of the surrounding height. Figure 4.44 shows an example of a large artefact corrupting the signal. In this case, the 'quality' level decreases periodically during the artefact.

4.6.3.1 Summary of Normalisation Effort

The 'normalisation effort' method works very well, even when the artefacts are short-lived (such as the one presented in Figure 4.43) and is relatively easy to convert to using streaming data. One major drawback, which is shared with the FFT heart-rate extraction method, is that the result does not remain in its 'bad state' (incorrect heart-rate versus poor 'quality') when there is an artefact; instead, it fluctuates depending on the signal. For the FFT method, a technique was developed to counter this error (described in Section 4.5.2.1), for the normalisation effort, a similar method could also work. Alternatively, the quality output could be expanded such that a single artefact would result in a lower quality value that spans a much greater window. This way, artefacts that have a behaviour like the one in Figure 4.44 (oscillating quality factor), would have their 'low quality regions' merge into each other creating a single large low-quality trough.

4.7 Combination of Methods

4.7.1 Heart-Rate

For heart-rate extraction, the three presented methods (peak-detection, FFT, and auto-correlation) could realistically be run in parallel, with their outputs combined in one of several ways. The simplest method would be to average the heart-rate results, which would *reduce* the error if one failed to obtain the 'correct' result. As each of the methods is potentially susceptible to different types of artefacts, then this could be an acceptable solution.

A second method that is often used in the digital world is the majority decision, this rejects values that do not align themselves with the majority of the other values. For three digital (binary) signals, the eight possible combinations of 0 and 1 can be grouped into two sets: the first with the

'ones' being the majority (011, 101, 110, 111), and the second with the 'zeros' being the majority (000, 001, 010, 100). For continuous signals, this technique cannot be applied directly as the signals are unlikely to completely agree with each other (have identical values). Therefore, a method must be used to either accept a pair of values (representing the majority) or reject a single value (the minority). In this case, rejecting the 'outlier' is appropriate and would be an improvement over the averaging method. The method in question, in this case, would reject the value that is furthest from the other two (the two values that are closest together are kept).

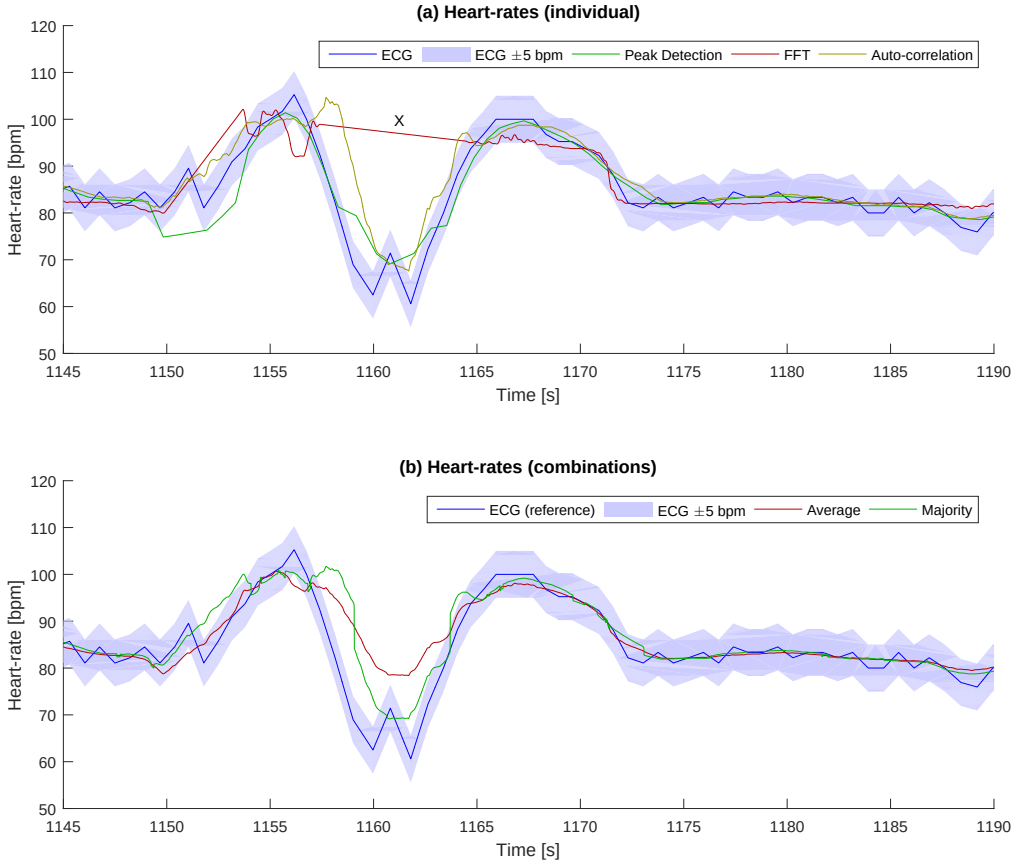

Figure 4.45: *The FFT method in the upper plot has failed to detect a heart-rate for this region (X), causing the post-processing algorithm to interpolate 'over the gap'. The other two methods (peak-detection and auto-correlation) closely agree on a heart-rate and so are chosen. The lower plot shows the two combination processes (averaging and majority).*

Examples of how effective the continuous 'majority' method is are illustrated in Figures 4.45 and 4.46. The first figure demonstrates one of the previously discussed errors with the FFT method, whereby a heart-rate is incorrectly attained, and the post-processing effectively 'draws a line' between the two nearest good heart-rates (linear interpolation, labelled 'X'). In this case, the other two heart-rates that follow the reference ECG value more closely are chosen and their average is used.

Figure 4.46 shows a similar case where one method's output is considered an outlier and removed from the average. This time the difference the rejection has made, has made the resulting average lie within the ±5 bpm window[16].

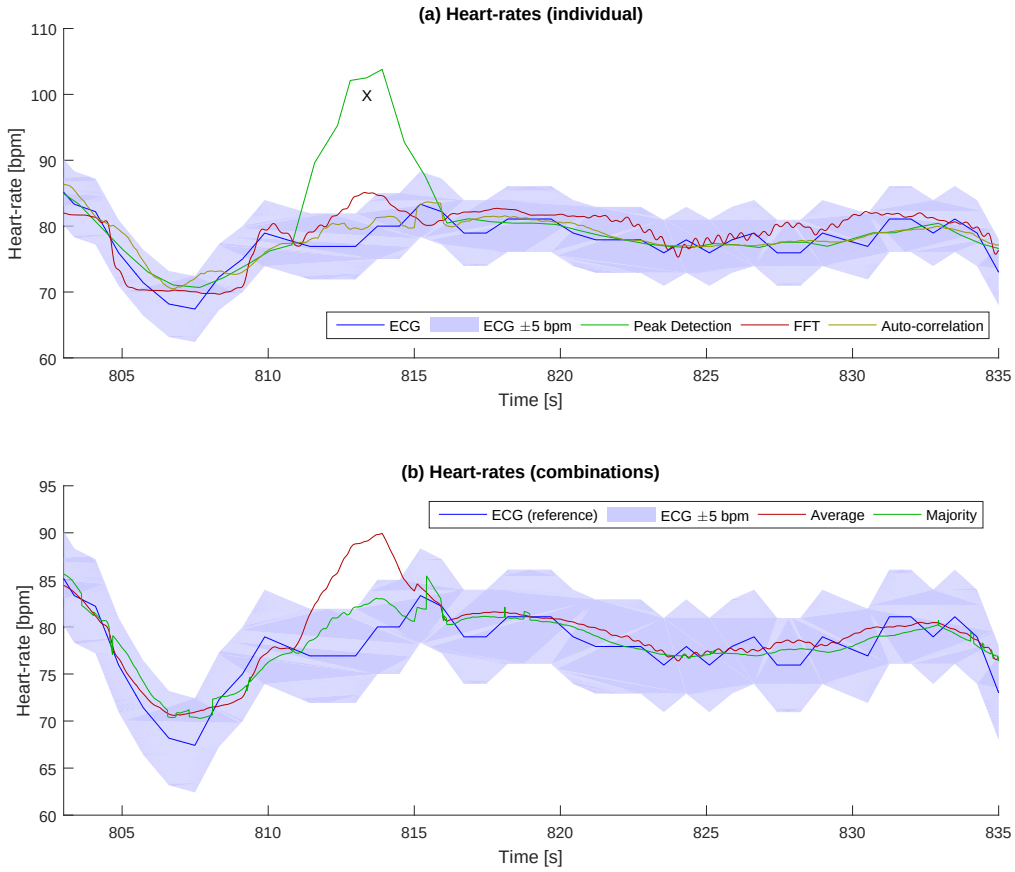

Figure 4.46: *The peak-detection algorithm (in the upper plot) has overestimated the heart-rate at this location (marked with an 'X') causing a large deviation from the reference. The lower plot shows the combination of methods (averaging and majority) of the heart-rates.*

[16]This window of ±5 bpm is arbitrary for these experiments, but was chosen to make an absolute target.

Figure 4.47 demonstrates one (potentially unavoidable) problem with the majority method: when the majority gets it *wrong*. In this example, at around 1115 seconds, the peak-detection algorithm performs the 'best' - with the minimal error from the reference - yet it is the one that is rejected due to the other two outputs (FFT and auto-correlation) being more similar, but incorrect. In this one case, averaging the three methods proved to be the best method (see lower plot in Figure 4.47).

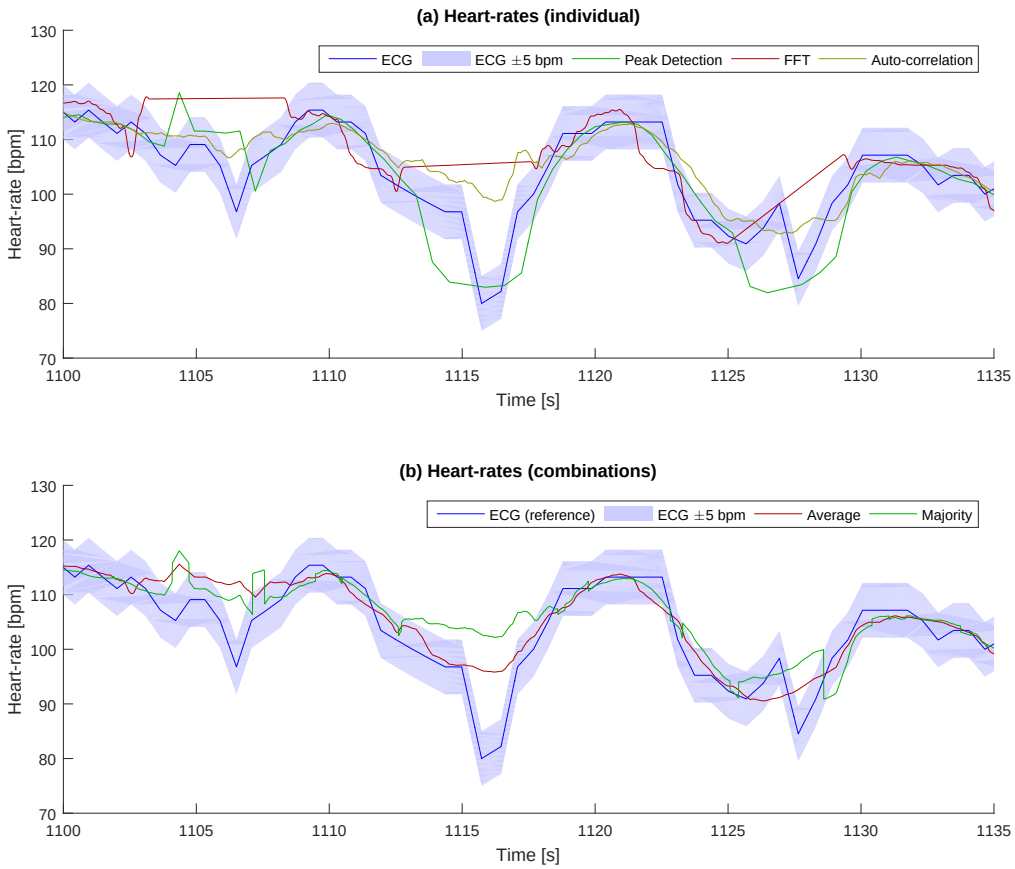

Figure 4.47: *An example of where the simple averaging technique outperforms the majority method. The lower plot shows the combination of methods (averaging and majority) of the heart-rates.*

4.7.2 Quality Measure (Indirect)

As the three outputs from the 'quality measure' methods (out-of-band analysis, envelope analysis and normalisation effort) do not give directly comparable values (like the absolute heart-rate values from the three heart-rate

extraction methods), there is no trivial way to combine them. Instead, they can be used individually to assess whether the heart-rate value(s) obtained are viable or not.

4.8 Discussion

Whilst it was not the intention to demonstrate a *complete* solution to determining the quality of a PPG signal (i.e: one that could be directly transplanted into a piece of equipment), the six algorithms presented in this chapter, of which three are able to also extract heart-rates, can be used to provide a starting point to such an algorithm. For the three methods that are used to assess the quality using attempted heart-rate extraction (peak-analysis, FFT and auto-correlation), the resulting heart-rate value can be directly used to reject 'bad' regions of the PPG waveform: regions where the heart-rate is unrealistically high or low. However, regions where there is a reasonable heart-rate output do not guarantee that the values are correct (since artefacts have the potential to 'create' PPG-like signals). This is the reason for developing the three indirect methods: out-of-band, envelope and normalisation analyses. These methods look at characteristics of the raw signal itself, and determine whether the signal is likely to contain an artefact or not.

Whilst these independent algorithms *could* be combined to form a single quality estimation function (and one for heart-rate extraction), the individual (indirect) methods will be used separately in the following chapter.

Chapter 5

In-Contact 'Matrix'

This section presents the reasoning behind; design and development of; and obtained results from, a piece of equipment that is used to extract PPGs from a subject, that also allows for motion artefacts to be suppressed when the resulting data is processed. The developed equipment is the final result of the research project and shows that the original goal (*motion artefact reduction for reflection-mode photoplethysmography*) can be achieved, at least in this set-up. Through the research performed in this project, several key findings have been gained, which have in-turn influenced both the path the research has taken, and the design of this final device.

For in-contact PPG measurements, relatively few 'types' of systems have been proposed and used. These systems most often take the form of one or two LEDs (of different wavelengths), a single photo-detector, an amplifier (often transimpedance), a filter, and an ADC[1]. Although other topologies exist (using timed charging, for instance, whereby the time for a photo-current source to charge a capacitor to a nominal value is used [90]), the simplicity of this design has made it prevail.

The intention for this research was *not* to experiment with other topologies, instead, in this chapter, the 'standard' topology as stated will be explored, but with a different physical configuration (and scale). This configuration was hypothesised to work after observation of the results from the previous two chapters.

[1]For digital systems. Some older systems will use an entirely analogue set-up.

5.1 Aims

The aim of this section of work was to design and create a novel system that could extract PPG signals from a subject with an increased tolerance to motion. After the system was designed and implemented, a participant trial could then be organised to collect data using the new equipment. The algorithms designed and tested in Chapter 4 would be used for the verification of the design.

5.1.1 Hypothesis

The non-contact experimentation presented in Chapter 3 provided, amongst other things, information regarding how the movement of the sensor *relative to* the subject affected the signal; and that minimising the velocity between the two was critical in maximising the 'quality' of the PPG.

The most obvious way to reduce relative, lateral, movement between the subject and sensor, was to move to an in-contact system so that the *imaging* sensor head was mechanically coupled to the subject, and was measuring the backscattered light directly from the surface. This would also eliminate any specularly reflected light as the source and sensors would occupy different spaces. The physical size of the imaging sensor would allow for detection of the PPG from multiple locations simultaneously.

However, by placing the sensor on the skin, any light that was used in the set-up would have to travel either from the 'far-side' of the subject (making it a transmission-mode system, restricting its use to limited locations on the body), or back-scattered beneath the surface from around the edge of the sensor (reflection-mode). The pixels in the centre of the image would 'see' a reduced amount of light. For an in-contact, reflection-mode, *imaging*, system that detects PPGs from a uniformly illuminated region, each pixel would have to have its own nearby light source. To achieve this,

a system could be created that behaves exactly like a 2D array (or *matrix*) of independent 'single-sensor' systems.

The hypothesis is that if multiple PPG sensors are in close proximity and placed on a subject, mechanically coupled to each other and the skin, then during motion the polarities of the PPG signals will remain constant, *but*, the polarities of the artefacts *may* not. If this is the case, then data processing can be applied to the resulting signals in order to extract the PPG signal (as desired) whilst suppressing the unwanted artefacts. Figure 5.1 illustrates the technique, whereby a movement that causes a compression of one 'side' of the sensor will have an opposite effect to the side that is 'released'. In theory, if this is the case, by simply taking an average of the signals, the amplitude of the differential 'artefact' can be reduced, whilst increasing the amplitude of the 'common mode' PPG. Additionally, by observing the relative qualities of the individual channels, some that are considered 'poor' can be excluded, increasing the overall quality.

At the time of designing the system, the easiest method of creating the 'sensor head' was to use a bundle of fibre cables that were packed in the desired configuration, which in turn connected to a processing unit that separated the LED sources and photo-diodes into larger, easier to design and build, banks. This way, the physical electronics do not need to be miniaturised in order to test the theory. If everything works as expected, then there is the potential to make a smaller, more 'user friendly' device (see Section 6.2).

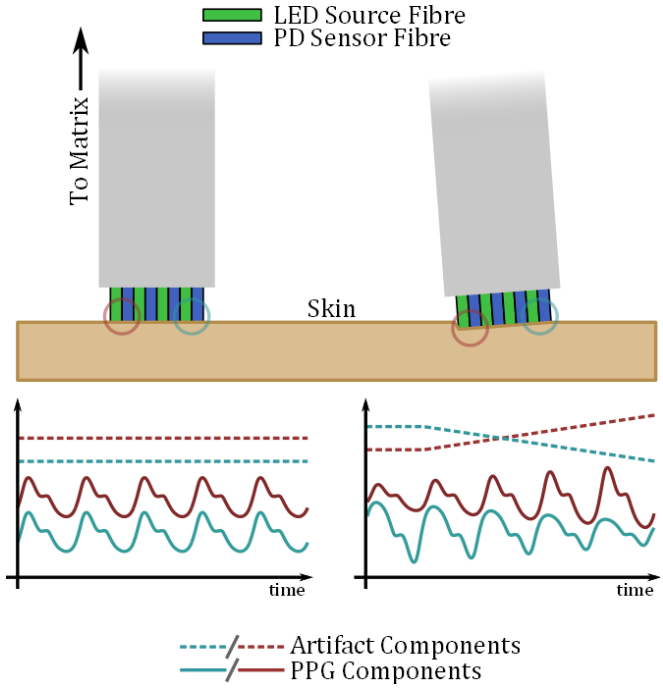

Figure 5.1: *Illustration of the idea that movement of a 'matrix' of individual sensors can cause predictable and measurable patterns in the resulting data, which can be used to compensate for the artefacts. In this simple case, the polarities of the 'artefacts' ('DC' component) are opposite, whilst the PPG polarity remains constant. Summing the two signals would 'cancel' the error.*

5.2 Design

In order to create a simple, reliable, multi-channel system, the decision was made to **not** make it portable. This was for two primary reasons; firstly, the amount of data that was capable of being generated would not have easily been 'streamed' to a PC, so it would have to have been stored. With the knowledge gained from designing and *using* the portable PPG/ECG recorder (see Section 4.2), although the data could be saved to a memory card for later processing, it was extremely desirable to see the data in real-time such that a fast iteration of experimentation and correction (trial and error) can be used. Secondly, by mounting the system on a desk, it can be made larger, allowing for more convenient components to be used (components that are easily hand-soldered, for example). This makes the design simpler, and reduces the construction cost and time. Additionally, a

pseudo-modular configuration can be used (that would be difficult if miniaturising the system), in order to make the system expandable, if needed. The 'modules' of the system also can be tested individually, then made in duplicate, further simplifying the construction.

The overall design (block diagram) is illustrated in Figure 5.2, and shows the duplicated blocks of light sources and sensors. The 'source' blocks each contain four green (525 nm) LEDs, controlled by a current source, fed from a single, four-channel, DAC. The 'sensor' blocks each contain four matching photodiodes and a transimpedance amplifier, multiplexed to a single ADC. All ADCs and DACs are controlled by an FPGA (simply due to there being four SPI ports for the ADCs, and four I^2C ports for the DACs, and micro-controllers with that many serial ports are rare). Whilst the I^2C and SPI ports could be shared (I^2C handles this directly, and SPI can be 'daisy-chained'), to get a higher data throughput, individual ports were used: not a problem in an FPGA where their respective controllers can be easily replicated in logic.

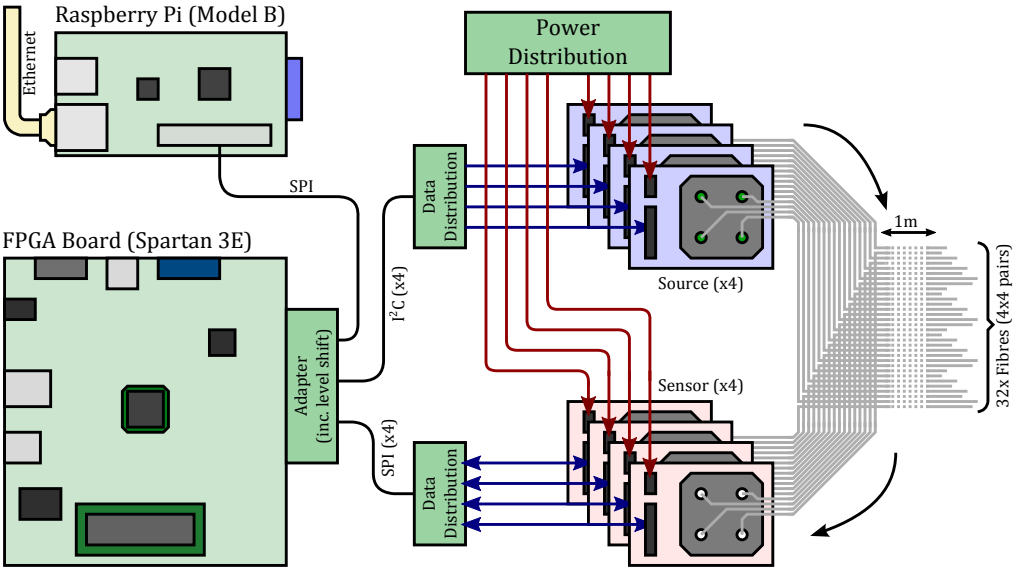

Figure 5.2: *Simplified block-diagram of the 'optical matrix' system. The 32 optical fibres (each with a 2 mm diameter) are joined to the PCBs via a custom (3D-printed) mount.*

Despite the FPGA development board (Spartan 3E) that was used con-

taining an Ethernet connection (and the necessary hardware to drive it), a Raspberry Pi was used to connect the FPGA to a PC, effectively behaving like an SPI to Ethernet converter. This was due to the relatively high development time of using the 'built-in' Ethernet connection, and the simplicity of using the Raspberry Pi. This also meant that, if needed, the Raspberry Pi could perform additional processing and buffering of the data.

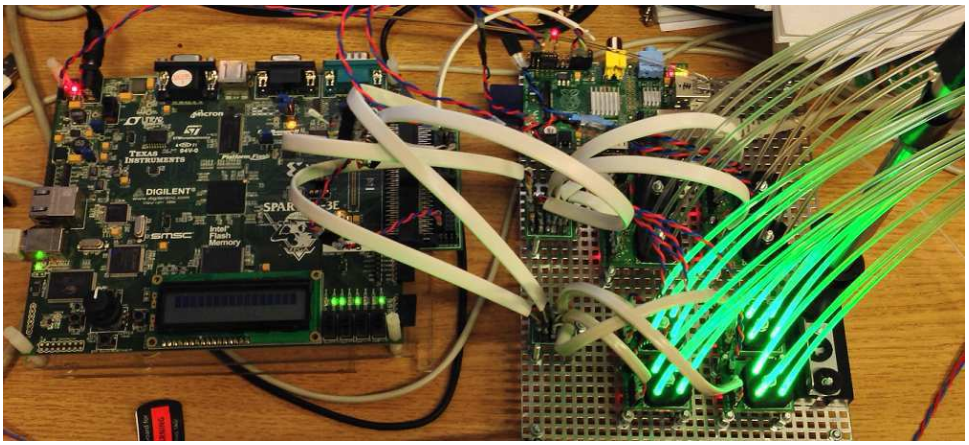

Figure 5.3: *Photograph of the complete system (with LEDs illuminated). The majority of the system's size comes from the FPGA development board, of which a small fraction is being used; i.e: the design has the potential to be minimised dramatically.*

Figure 5.3 shows the completed system (the schematic and PCB layout for the FPGA development board 'add-on' are shown in Appendix III). Three different fibres were tested during construction, each with a different diameter. A compromise had to be made between smaller, more flexible fibres, and larger fibres that allowed a greater amount of light to be transported. In the end, 2 mm (diameter) fibres were chosen. To minimise the cost of the system, and to simplify the construction of the 'sensor end' of the bundle, unsheathed fibre was used (part number `OMPF2000` from 'Fibre Data'). For the data collection using this equipment, the office lighting was switched off.

To interface with the hardware, a simple application was written to view the data in real-time, control the pattern/intensity of the LEDs, and save the data to a file for later processing. This application is shown in

Figure 5.4, and depicts a recording that was being made of a participant/-subject.

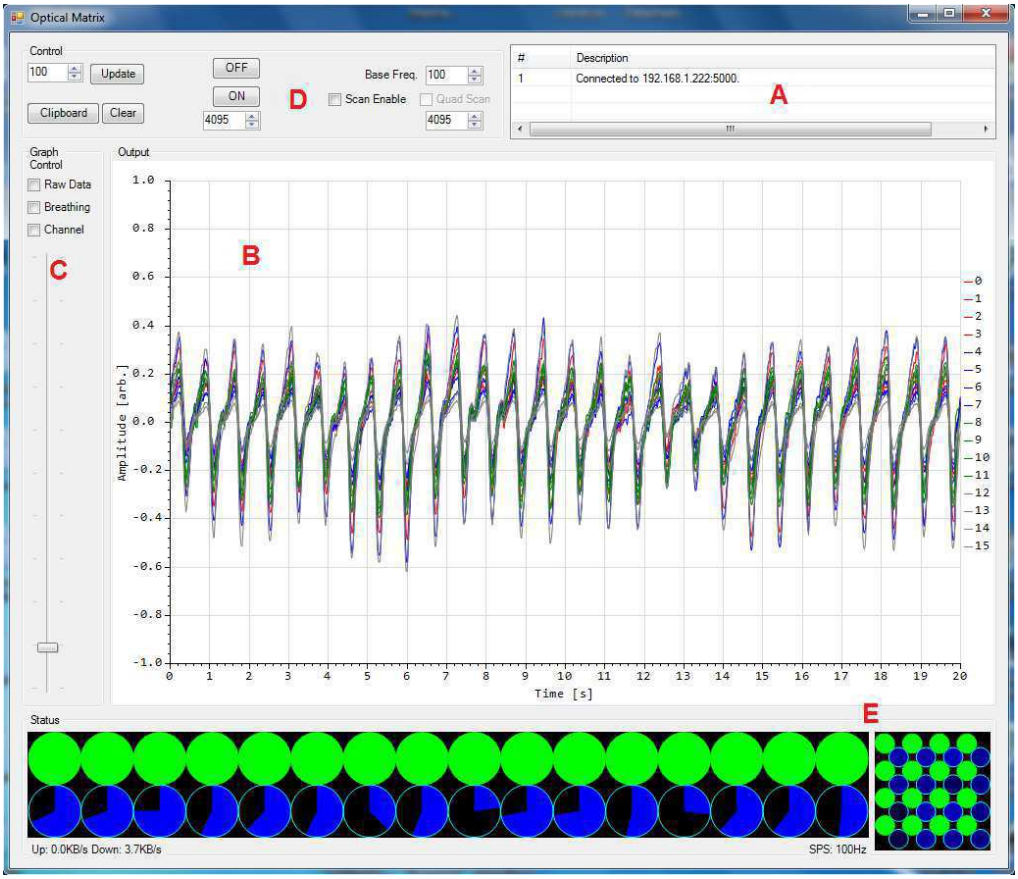

Figure 5.4: *Screen-shot of the 'optical matrix' interface. A: Hardware information (IP Address, etc). B: Output waveforms. C: Display controls (raw vs. filtered, etc). D: Output controls and file saving (LED intensity, and frequencies, etc). E: Visualisation (pie-chart) of the status of each of the channels (output and input illumination).*

The simplest 'mode' the optical matrix can be configured in is to simply illuminate all LEDs with a fixed intensity, and for all 16 of the sensors to be sampled and their data output at a constant rate (typically, 100 sps was chosen, to mimic the data-rate from the PPG/ECG recorder). This would then be transmitted to the PC for viewing, and saved to a file, where it could be processed at a later date. Additionally, different patterns of illumination can be programmed; one example of an illumination pattern is for each LED to be illuminated in turn, whilst still capturing all 16 sensor channels per iteration. Another pattern splits the 16 sources into

four quadrants, with some receiving channels being in locations where the light is not present (or at least having less of an influence).

5.3 Protocol

Participants in a data-collection trial were asked to place their palm (specifically, the Thenar region, below the thumb, see Figure 5.5) on the fibre-bundle sensor head, and perform a series of 'activities' in order. These consisted of holding still, rotating their hand in each of the three axes (A, B, and C, in the figure), and finally, 'compressing' their hand into the sensor (D). The study was approved by the University of Nottingham's Research Ethics Committee.

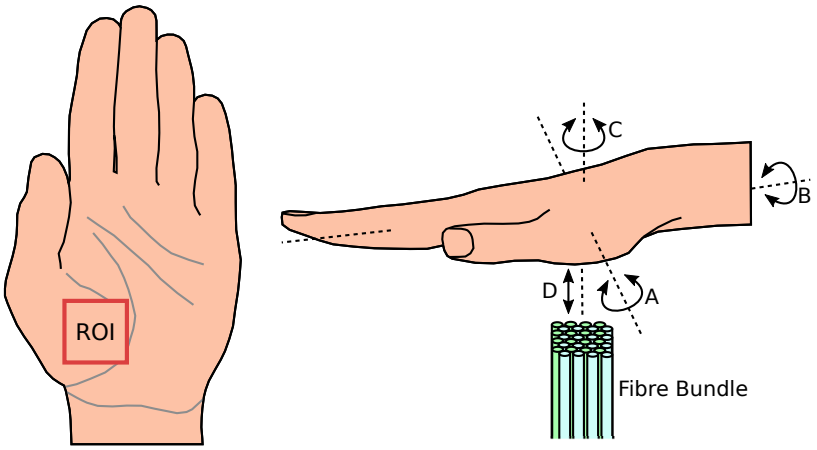

Figure 5.5: *Illustration of how the fibre bundle (the sensor head) was positioned on the subject's hand, and the movements the subject was instructed to make.*

For the main part of the trial, the participants that volunteered (two, both male, white) were asked to, between holding still, move their hand (that was placed on the end of the fibre-bundle) in different directions. These directions were:

Compression The hand was moved in the same direction as the fibres, such that each fibre pair 'saw' the same motion.

Rotation 'X' The bundle was rotated about the end-point, such that one 'side' of the fibre pairs were pressed more firmly onto the skin, whilst the other side was released.

Rotation 'Y' The bundle was rotated about the end-point, as above, but in the orthogonal direction.

Twist The bundle was rotated in the same axis as the fibres such as to lightly stretch the skin. Note: the ends of the fibres did not 'slip' over the surface.

The participants were asked to do each movement for at least five seconds, with a similar period of inactivity in-between.

5.4 Results & Processing

To verify the effectiveness of the processing techniques that are discussed, the three 'indirect' quality measurement algorithms were run on the recorded data, and the periods of movement activity compared to the regions of inactivity.

Originally, the scanning functionality of the Optical Matrix was intended to be used, but the results did not turn out as expected as described in Section 5.4.1.1, therefore the scanning data was not processed. The non-scanned data was used instead.

5.4.1 Raw Results

Data from both participants is presented in Figures 5.6 and 5.7. For clarity, the signals have been band-pass filtered with the previously used and described Butterworth filter (0.66 - 9.00 Hz), to remove each channel's drifting base-lines.

As the participants were being supervised during the recording, the times when the movements occurred could be recorded accurately. These regions are highlighted in the figures in red. There are four such regions in each recording, and the effect of the motions are clearly visible on each channel.

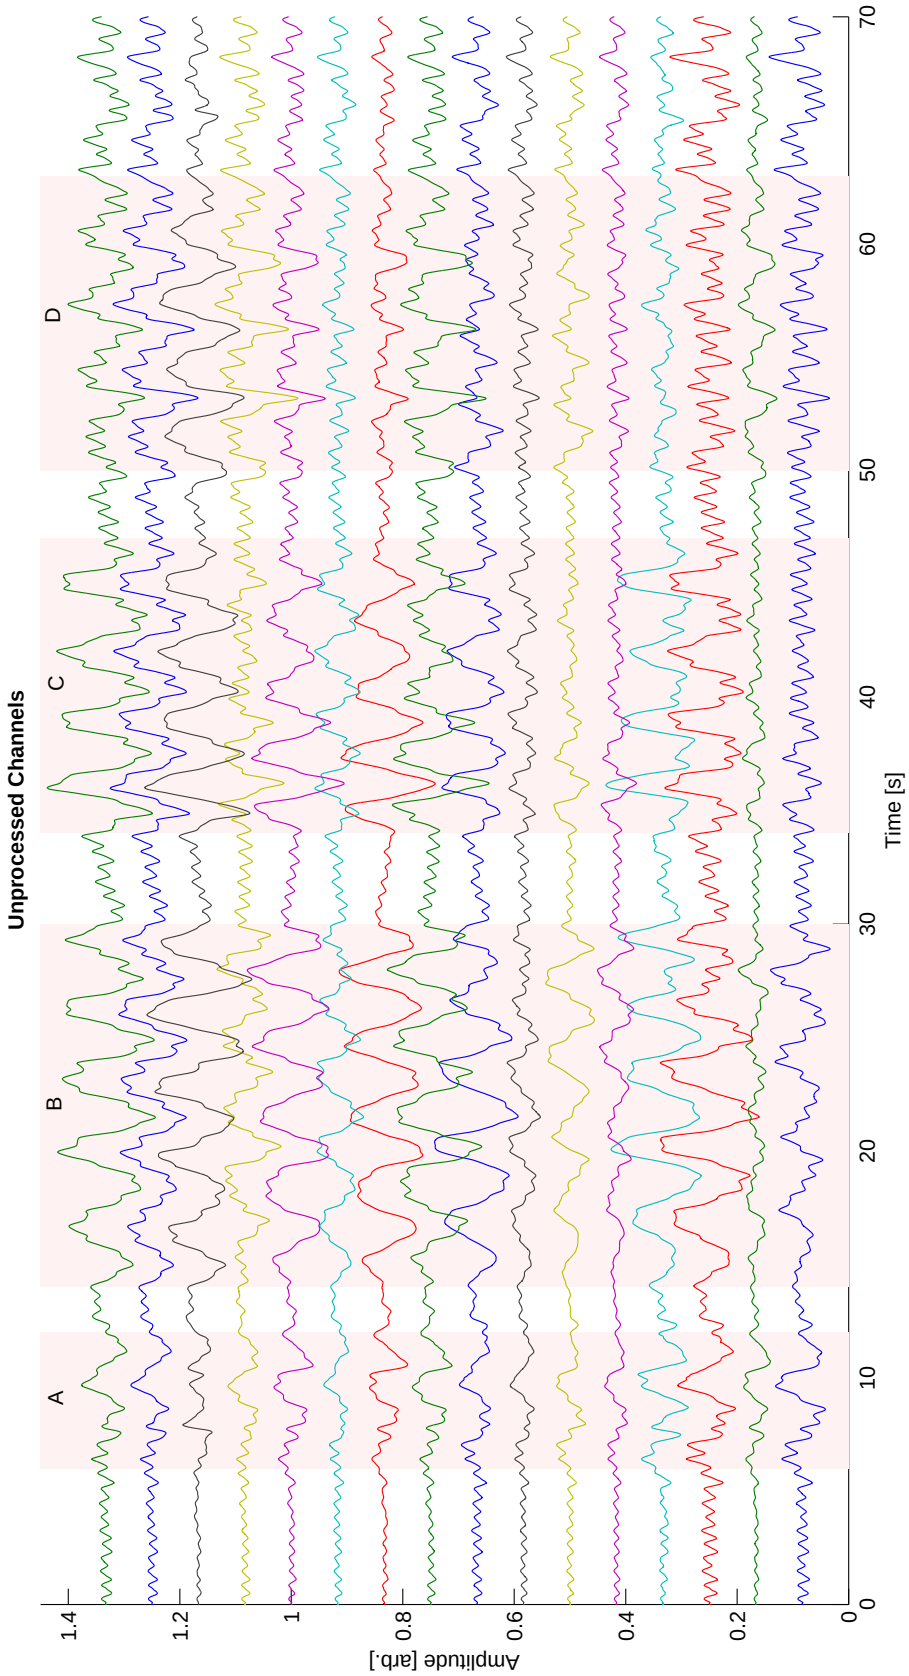

Figure 5.6: *All sixteen channels of participant 1's recording. All LEDs are illuminated equally with the same intensity.*

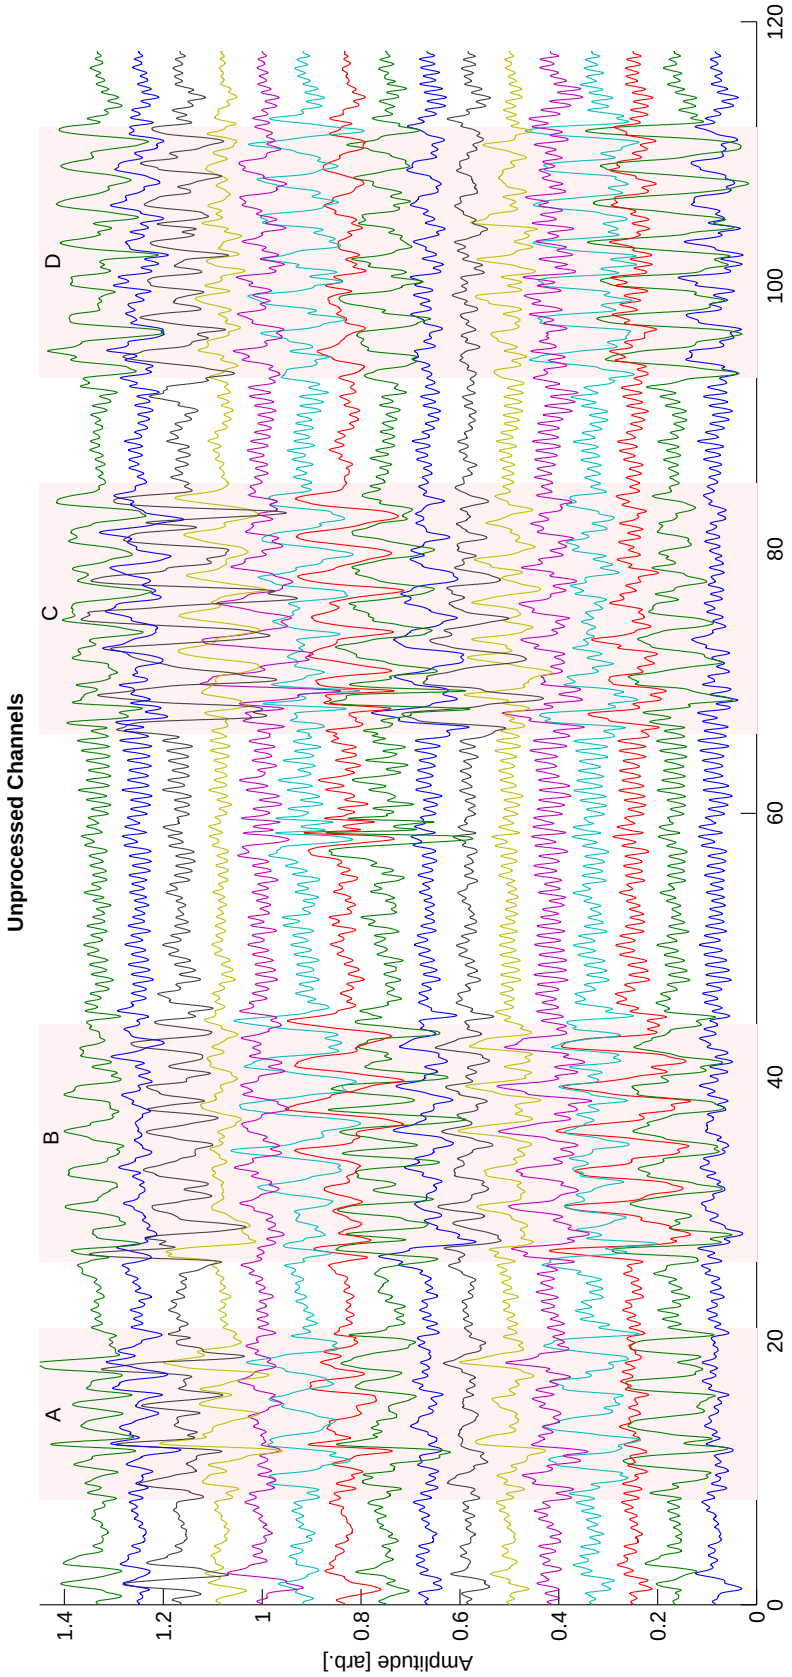

Figure 5.7: *All sixteen channels of participant 2's recording. All LEDs are illuminated equally with the same intensity.*

5.4.1.1 Additional Data

In addition to the recordings presented above, a different set of recordings were made that allowed for testing the functionality of the 'scanning' modes of the hardware. The sixteen fibre channels were separated into four 'quadrants', and the LEDs were illuminated in each quadrant in turn, whilst the data from all sixteen photodiode channels were saved for each iteration. In theory, the receiving channels that were close to the illuminating channels would get a strong PPG signal, whilst the channels that were further away would get a weaker one (or none at all).

Figures 5.8 and 5.9 show the resulting sixteen channels of data when just single quadrants are illuminated. When quadrant 4 (Figure 5.8), in the 'lower right' of the bundle, is illuminated (green circles represent the LEDs), some of the neighbouring detection channels 'see' the PPG (blue circles represent the photodiodes). In fact, for this iteration, there is at least one photodiode channel in each quadrant that is adjacent to an illuminated fibre, resulting in a strong PPG signal. Similarly, when quadrant 1 (Figure 5.9), in the 'upper left' of the bundle, is illuminated, due to the relative positions of the source and sensor fibres, only one receiving quadrant has the PPG present strongly (relative to the other channels).

Whilst it is possible to illuminate just a single LED to extract PPGs (see Figure 5.10), the fast fall-off of the light only allows for detection of a strong PPG in directly adjacent fibres. It was the intention originally to illuminate each individual LED in-turn, and look at the signals that arise from all detectors. In theory, by doing this, more detection sites could be 'created' by looking at non-adjacent pairs, hence increasing the resolution of the array. However, as the sensitivity of the system was not sufficient, this was not possible.

One possible alternative that may have allowed for non-adjacent sensors to pick up the PPG would be to move the fibres closer together (i.e: by

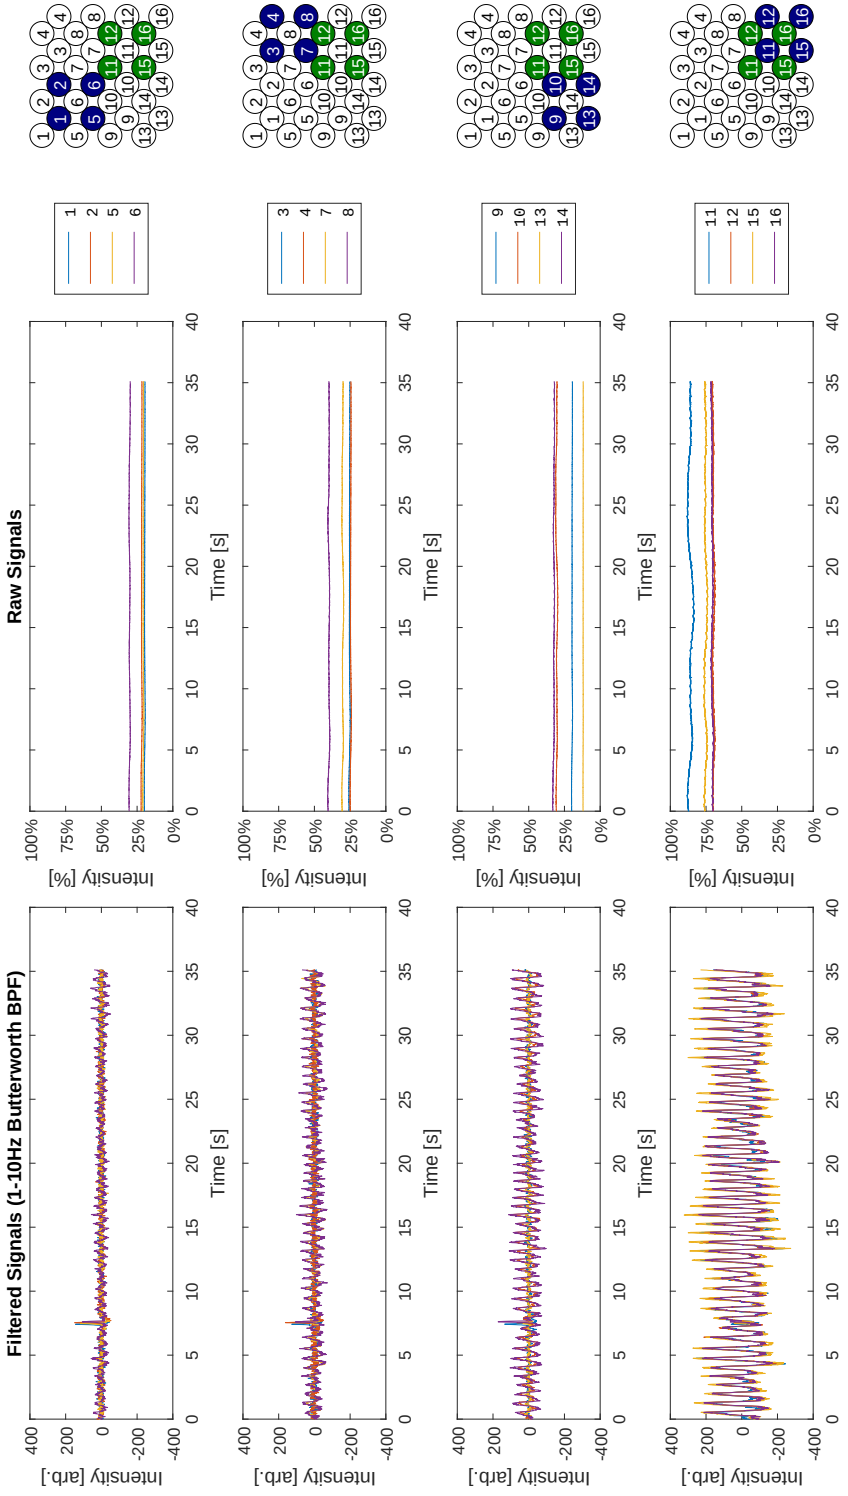

Figure 5.8: *Photodiode responses when just one quadrant is illuminated. The lower right quadrant (quadrant 45) is illuminated. Due to the neighbouring fibres, a PPG is being detected strongly in quadrants that are not being directly illuminated.*

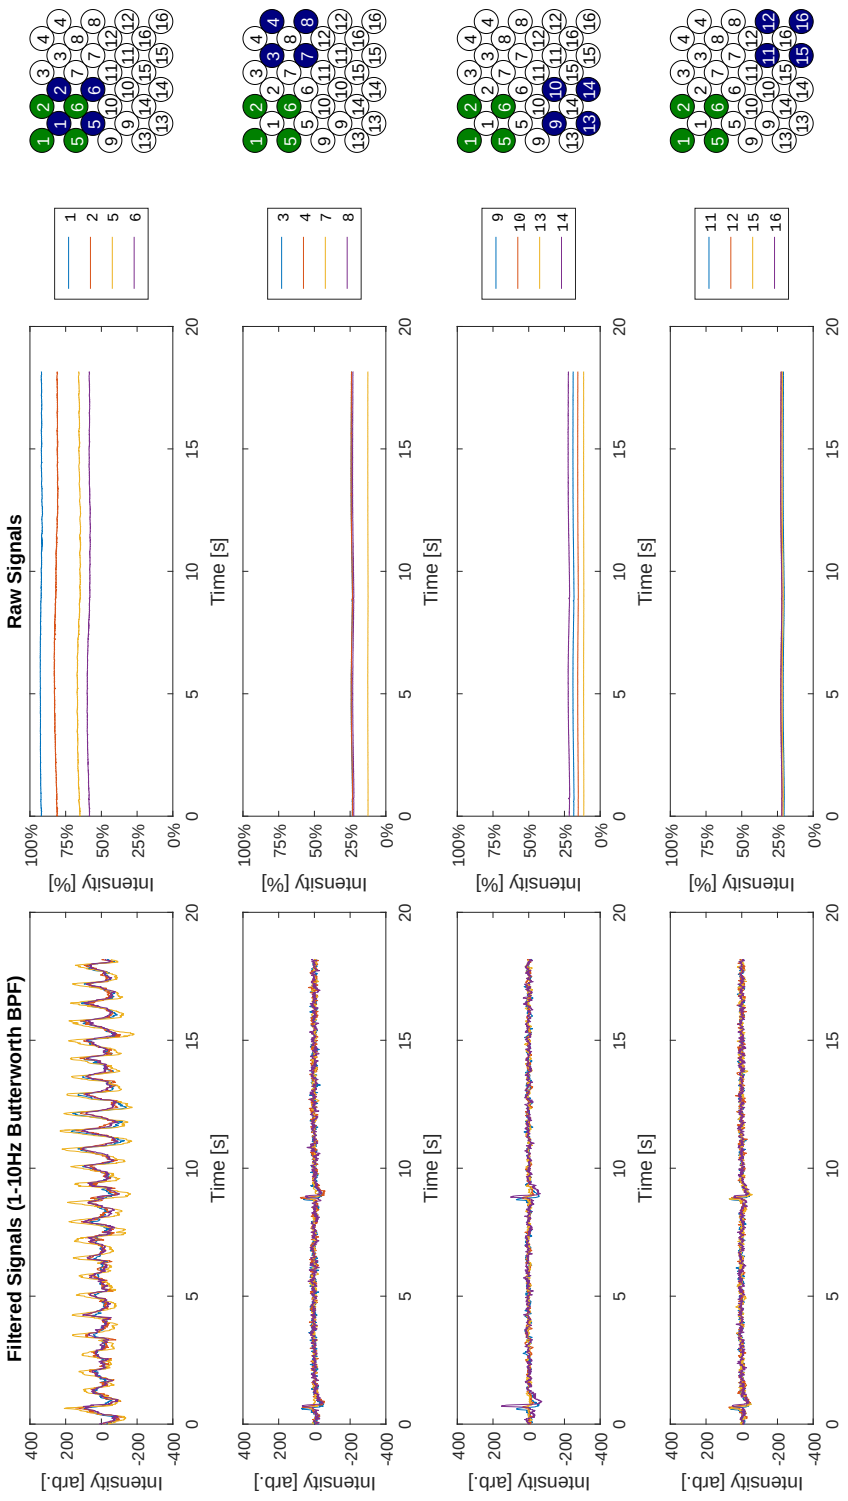

Figure 5.9: *A similar set of responses to the previous figure, but with a different quadrant (quadrant 1) illuminated. Since sensor fibres in the neighbouring quadrants are not directly adjacent to the sources, there is a much weaker PPG present in them.*

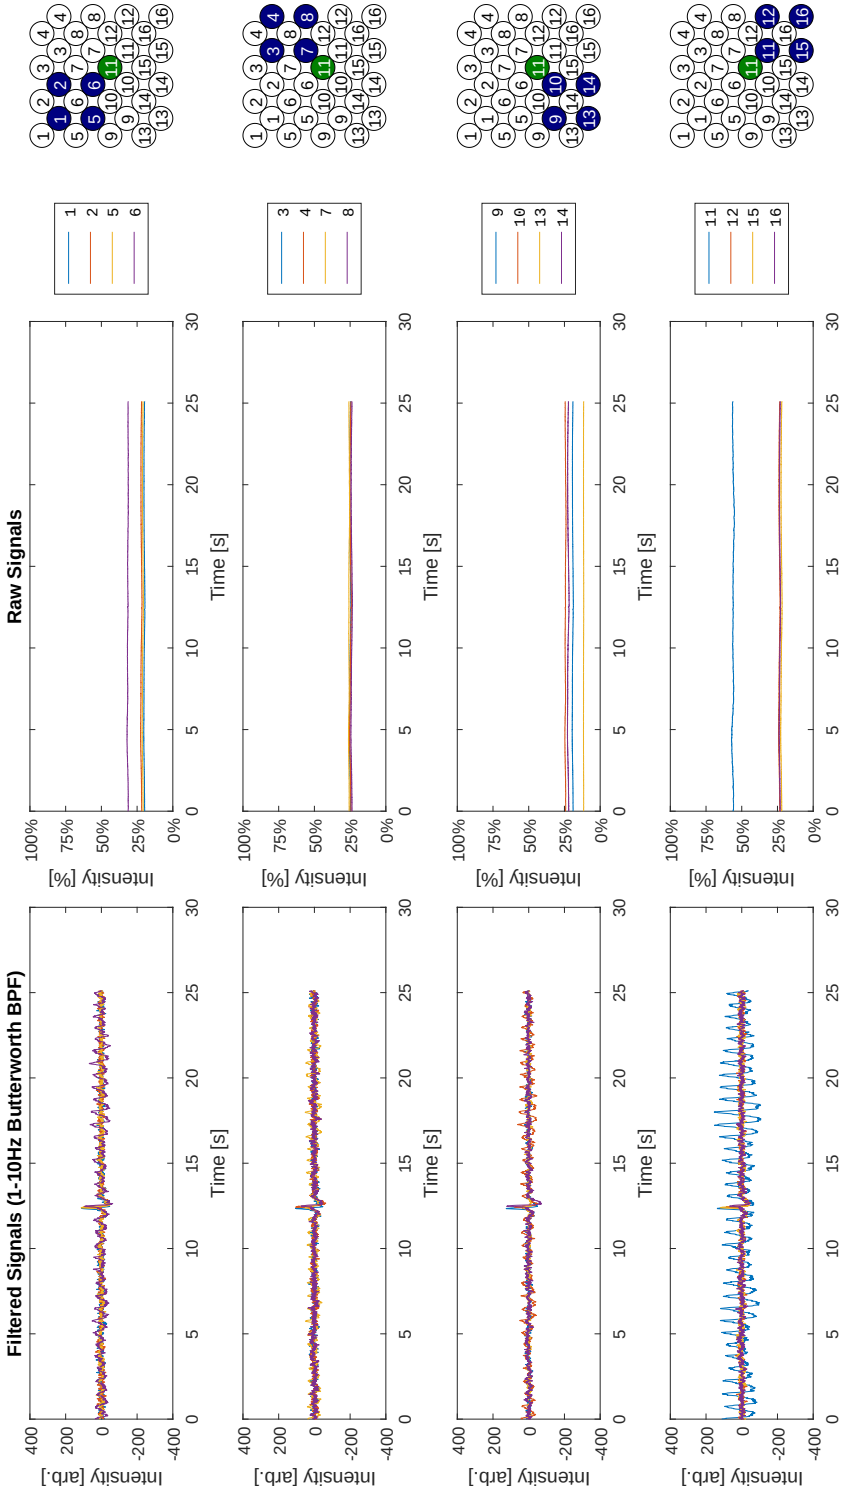

Figure 5.10: *Only the directly adjacent fibres are able to 'see' a strong pulsatile signal. The other fibres are too far away for the received data to be useful. With a more sensitive detector, or by using a higher light intensity, it may have been possible, however.*

using narrower fibre-optic cables). If shorter lengths were used to reduce losses, and better coupling was used at the interface between the cables and the LEDs and photodiodes, this may have worked.

Due to these limitations of the system, the intended use (of the 'scanning' modes) was not used. However, an alternate advantageous use, with sensors picking up alternating signals of illumination by an LED and effectively ambient light, then simple 'ambient light subtraction' as explained previously can be used. This allows for the design (with the unsheathed fibres) to be used without the need for turning off artificial lighting or blocking natural light[2].

5.4.2 Averaging

For *non-contact* measurements with a camera, it has been shown that the area of the region-of-interest has a direct effect on the 'quality' (the signal-to-noise ratio) of the obtained signal [69]. In general, by increasing the area, the uncorrelated components in different pixels (noise, etc), will be reduced, with the correlated signals (the PPG, etc) enhanced. As the optical matrix that has been created is, effectively, an *in-contact camera*, the same technique can apply.

By averaging all sixteen in-contact channels, the correlated components will be left alone, whilst the uncorrelated components will be attenuated. Figures 5.11 and 5.12 show this effect. Individually, in each of the sixteen channels, the previously described artefacts can be clearly seen (marked as 'A', 'B', 'C' and 'D' in Figures 5.6 and 5.7). But when averaged, all but one of the artefacts are attenuated.

The three movement types (compression, rotation, and twisting) that were tested each show different responses. At around 10 seconds into the participant 1's recording (Figure 5.11, A), for example, the fibre was pressed

[2]Artificial lights were turned off for the participant trials, anyway, however.

into the hand more firmly, and then released, for two cycles. As each individual fibre 'sees' the same artefact (i.e: the artefact is correlated between each sensor), the erroneous signal is still strongly present in the average. This is expected, simply due to the fact that in this scenario the fibre bundle behaves like one larger fibre - there is no variation between the individual fibres.

The two movements centred at 25 and 40 seconds (B and C), created by rotating the end of the fibre on the skin (such that one side of the fibres press harder, whilst the other side press more lightly), show the largest improvement. This can be attributed to the effective *differential* signal that is being created, and then removed due to the averaging as explained above.

The fourth movement, at 60 seconds (D), shows a moderate improvement over the individual channels. The artefact at this location was caused by the fibre bundle being 'twisted' on the hand, but held in place to prevent the end of the sensor 'slipping' over the skin. In the individual outputs, a common-mode signal is present in some of the channels (visible as large negative-going spikes between 50 and 60 seconds), but not in others. This is most likely due to the fibres in the centre of the bundle not moving as far and producing a smaller artefact. This is particularly evident in participant 1's raw data (Figures 5.6)[3].

[3]The 'centre of rotation' of the bundle for the participants are not located at the 'geometric centre', but is instead offset closer to an edge. This is due to the practical difficulties of aligning the fibre correctly during the trials.

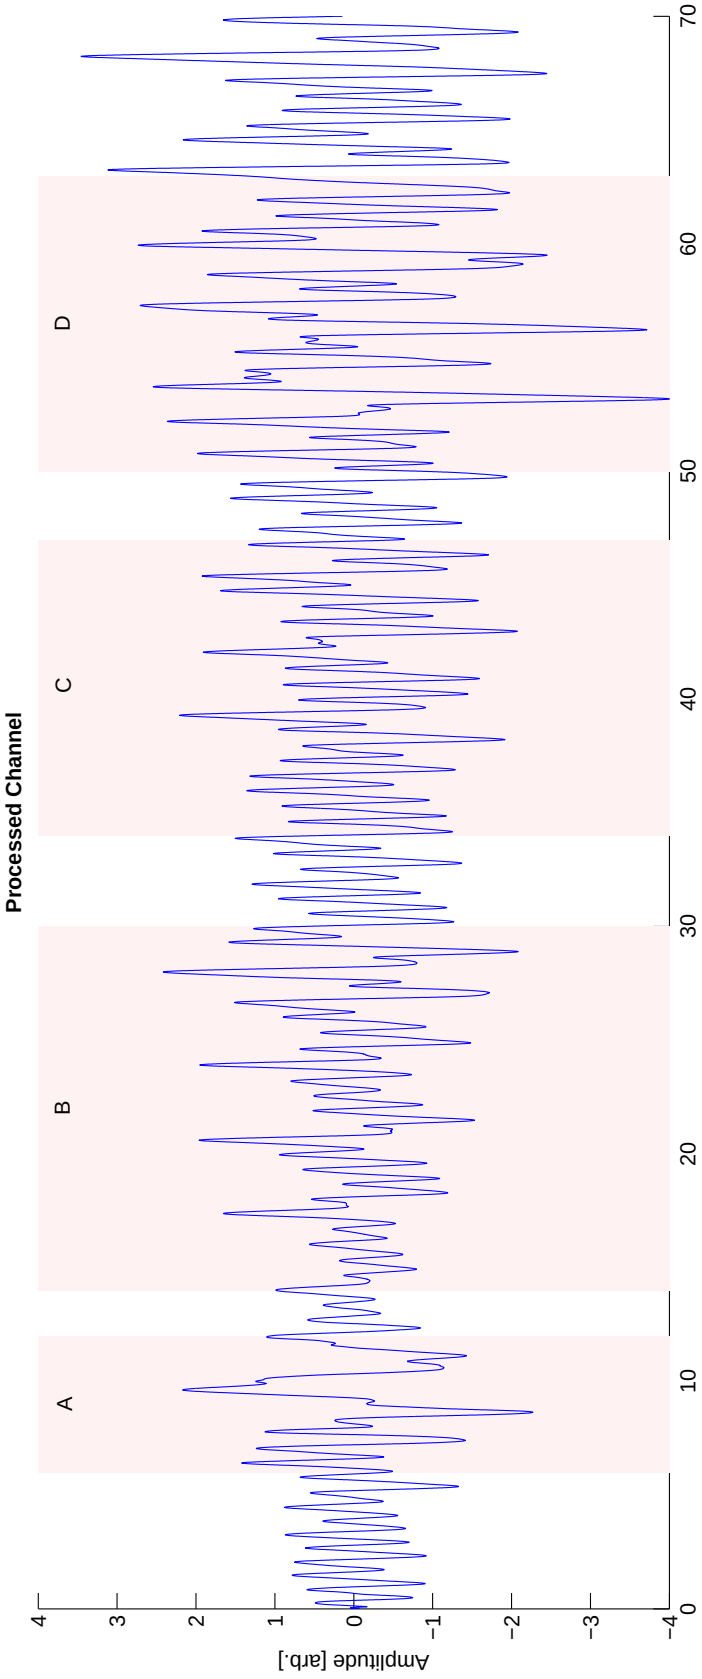

Figure 5.11: *Average of all sixteen channels from participant 1's recording. Some artefacts that are present in the individual channels appear to have been suppressed in the average, suggesting that they appear in different channels with differing polarities, as expected. All LEDs are illuminated equally with the same intensity.*

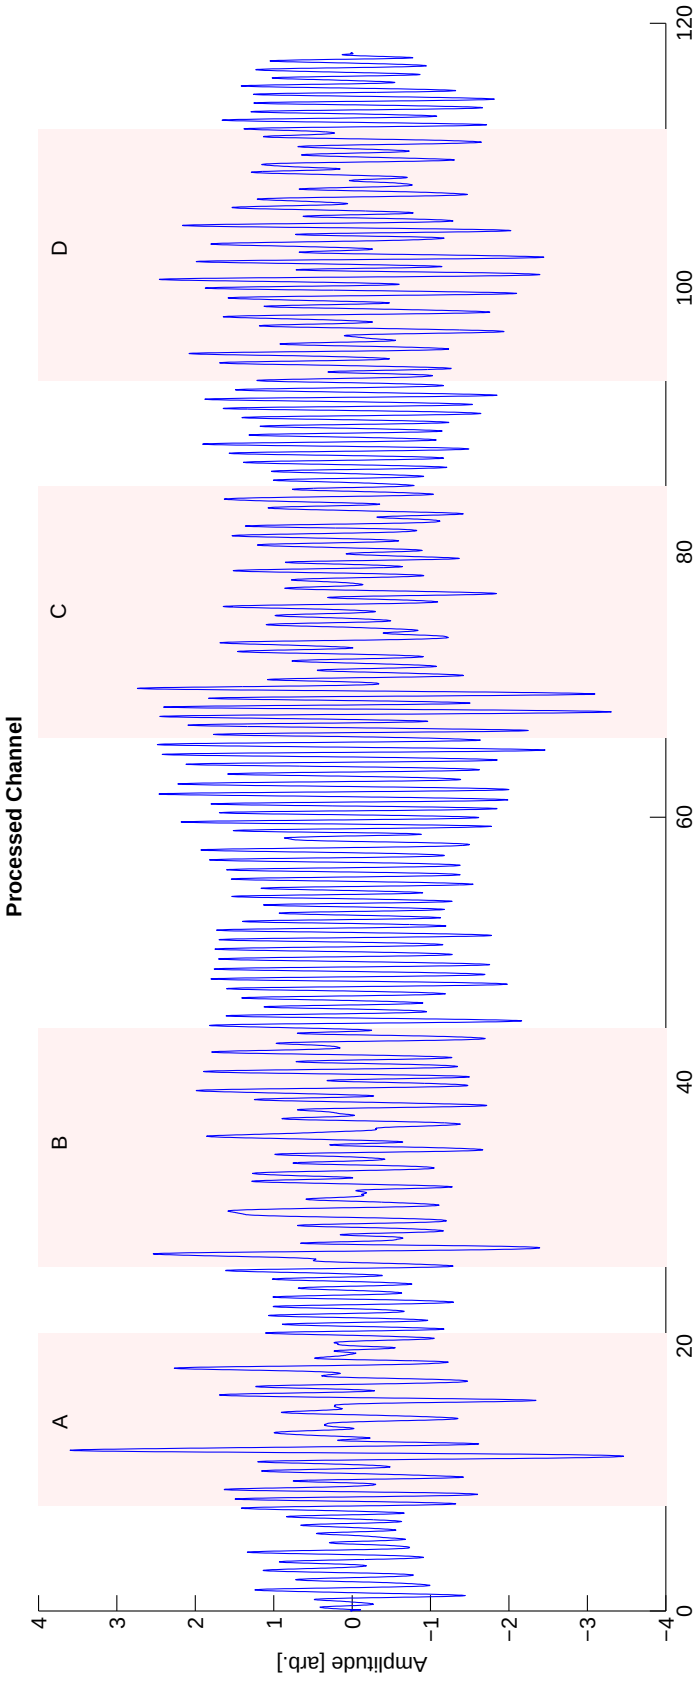

Figure 5.12: *Average of all sixteen channels from participant 2's recording. Some artefacts that are present in the individual channels appear to have been suppressed in the average, suggesting that they appear in different channels with differing polarities, as expected. All LEDs are illuminated equally with the same intensity.*

To give a better illustration of how the artefacts that arise from the rotational movements are suppressed, two waveforms from the sixteen channels (of participant 1) that straddled one of the 'axes of rotation' were selected for a direct comparison. These two channels are presented in Figure 5.13 and clearly show complimentary (differential) waveforms. Similarly to the 'twist' motion, the axis of rotation for this motion appears to be offset from the centre.

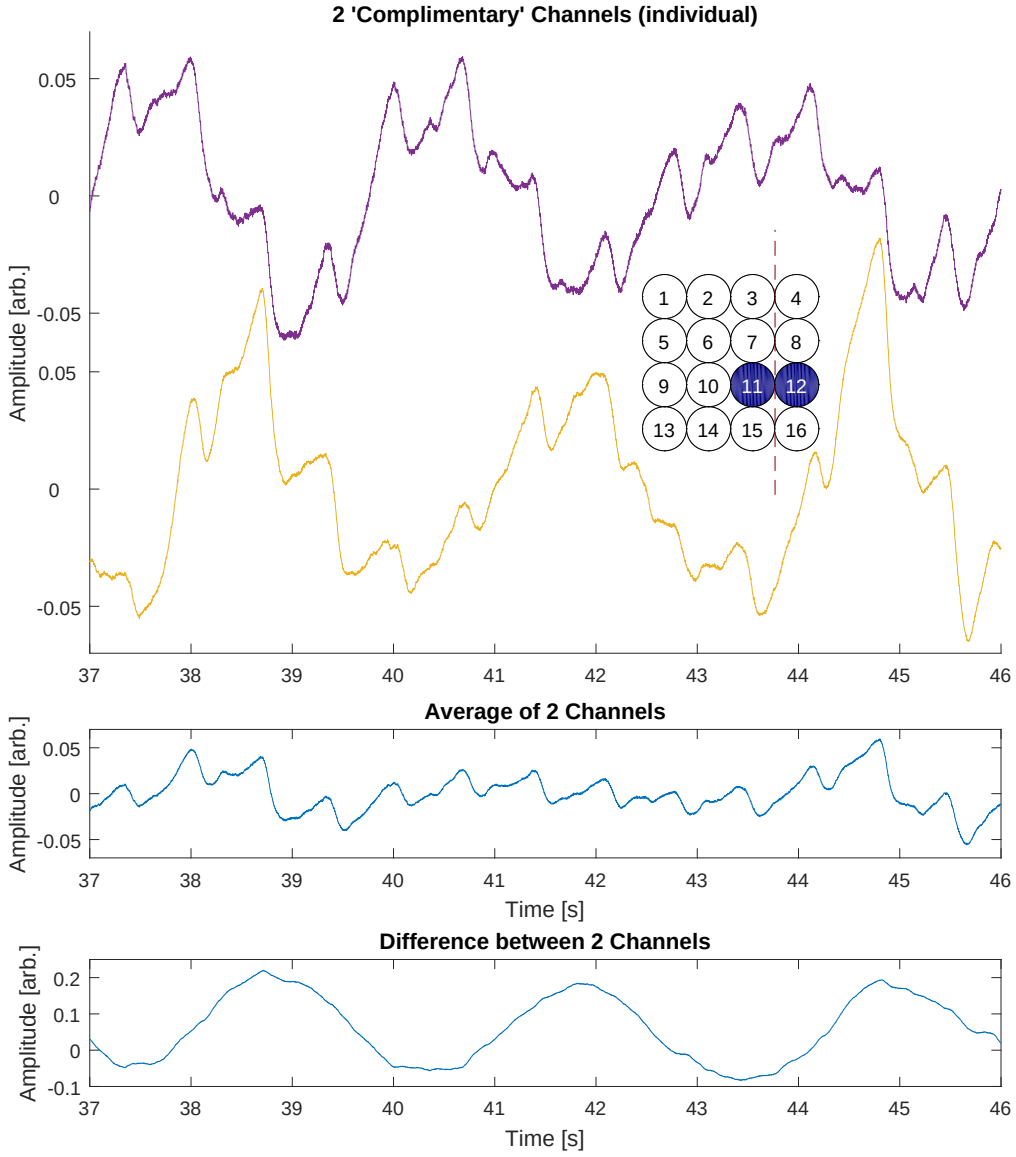

Figure 5.13: *Two complimentary waveforms from either side of an axis of rotation, along with their sum (average) and difference. The sum and difference of the signals effectively extract the differential and common-mode components, which in this case, are the PPG and the artefact.*

It should be noted that despite the axis in Figure 5.13 (the red dotted line on the illustration of the fibres) being portrayed as running 'vertically', it is in-fact angled away from the centre as well. The two channels, 3 and 4, for example, show a common-mode artefact, simply due to them both being on the same side of the axis.

As a direct comparison to the 'complimentary' channels of Figure 5.13, Figure 5.14 shows a pair of similar signals that have the motion artefacts with the same polarity. Here, the channel that is farthest away from the axis of rotation (channel 9) has a larger amplitude signal due to the end of the fibre moving over a larger distance during the rotational movement. The two processed signals (sum and difference) shows much degraded PPG waveforms[4].

As the axes of rotation (for both the 'rotation' movements, and the 'twisting' movement) do not divide the sixteen channels evenly, there is a slight imbalance in the differential-mode signals. This results in the differential artefact signals not being completely removed. If the axes were positioned centrally, then the averaging technique would be sufficient. However, as the pseudo-random movement is unlikely to align itself as desired, an alternative method of reducing the differential signals is required; one that is insensitive to the aforementioned imbalances (such a method is discussed in Section 5.4.3).

[4]Despite the artefact, in this case, being of a sufficiently different frequency to the PPG (such that it could be filtered out), the example demonstrates the principle.

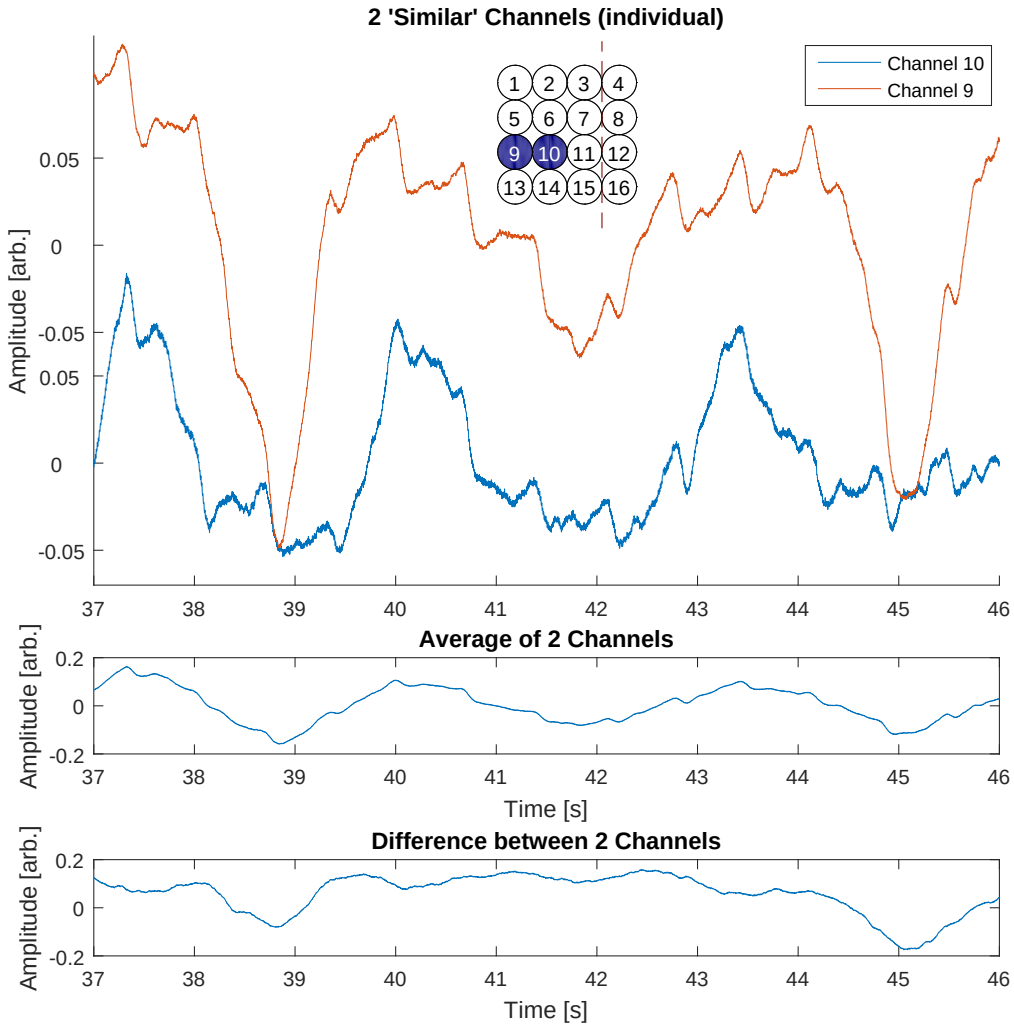

Figure 5.14: *Two similar waveforms from the same side of the rotational axis, along with their sum (average) and difference. As channel '9' is further from the axis, it will experience a larger deviation in the signal (the movement has a greater amplitude, mirrored in the resulting signal).*

5.4.2.1 'Quality' Algorithm Results

By passing the two participants' raw data through the previously developed 'quality calculating' algorithms, and comparing the results with the results of the processed data (average of the sixteen channels in this case), the effectiveness of the processing technique can be quantised.

Figures 5.15 and 5.16 show the results of the raw (16 channels) and the processed (averaged) data after run through the three 'indirect' quality measuring functions: 'out-of-bounds', 'envelope', and 'normalisation effort' analysis.

To simplify the graphs, of the sixteen channels the algorithms were run on, the *worst*, *average*, and *best* 'qualities' are shown. These are presented in blue, green, and red, respectively. The result from the processed data is shown in black.

In both participants' recordings, the regions where no movement was present shows a significant 'drop' in the 'anti-qualities'. The algorithms represented in the upper two graphs show a larger number for more 'chaotic' (lower quality) signals. The algorithm for the lower graph normally represents a true 'quality' (from 0 to 1), but has been inverted so as to keep consistent with the other graphs.

Statistically, none of the quality vectors from any of the three algorithms are normally distributed, as tested using the One-sample Kolmogorov-Smirnov test (with $p < 0.05$) [91]. This means that a valid comparison between the 'best channel' quality and the processed 'averaged' channels could be made using a Wilcoxon signed rank test [92]. For the 'out-of-bounds' and 'envelope' algorithms, the averaging of the sixteen channels yields a statistically different quality result ($p < 0.05$) than the best-case channel.

Tables 5.1 and 5.2 show the times-improvements of the averaging of the channels, over the best-channel quality results, for the 'out-of-bounds' and 'envelope' analysis algorithms, respectively. The 'normalisation' analysis algorithm cannot be included in these statistics tests as more than 50% (average between the participants) of the resulting quality data is identical, and thus cannot be declared as statistically different using the described method. Table 5.3 shows the percentage of time each algorithm outperforms the other.

Over the two participants, understanding that a quality level of 0 represents a perfect PPG, with a positive value representing a degraded signal, the averaging algorithm produces - on average between the participants

- a 2.5× improvement in measured quality (using a combination of the out-of-bounds and envelope methods).

Participant	**Average** *Median (IQR)*	**Best Channel** *Median (IQR)*	**Improvement** *(best/average)*
1	1.042 (0.935)	1.321 (2.135)	1.268 ×
2	0.373 (0.688)	1.363 (1.715)	3.658 ×

Table 5.1: *'Out-of-bounds' algorithm output statistics, compared to the 'best-case' channel, for both participants. All medians/IQRs* $\times 10^{-1}$.

Participant	**Average** *Median (IQR)*	**Best Channel** *Median (IQR)*	**Improvement** *(best/average)*
1	0.054 (0.132)	0.098 (0.158)	1.805 ×
2	0.300 (1.208)	0.978 (1.422)	3.257 ×

Table 5.2: *'Envelope' algorithm output statistics, compared to the 'best-case' channel, for both participants. All medians/IQRs* $\times 10^{-5}$.

Participant	**A >B** *(%)*	**A = B** *(%)*	**A <B** *(%)*
1	39.9	56.1	4.0
2	30.1	49.3	20.6

Table 5.3: *'Normalisation' algorithm output statistics, compared to the 'best-case' channel, for both participants.* ***A****: Averaging algorithm 'quality',* ***B****: Best channel quality. Where higher 'quality' is better (for example, 'A >B' represents the averaging algorithm outperforming the best-channel result).*

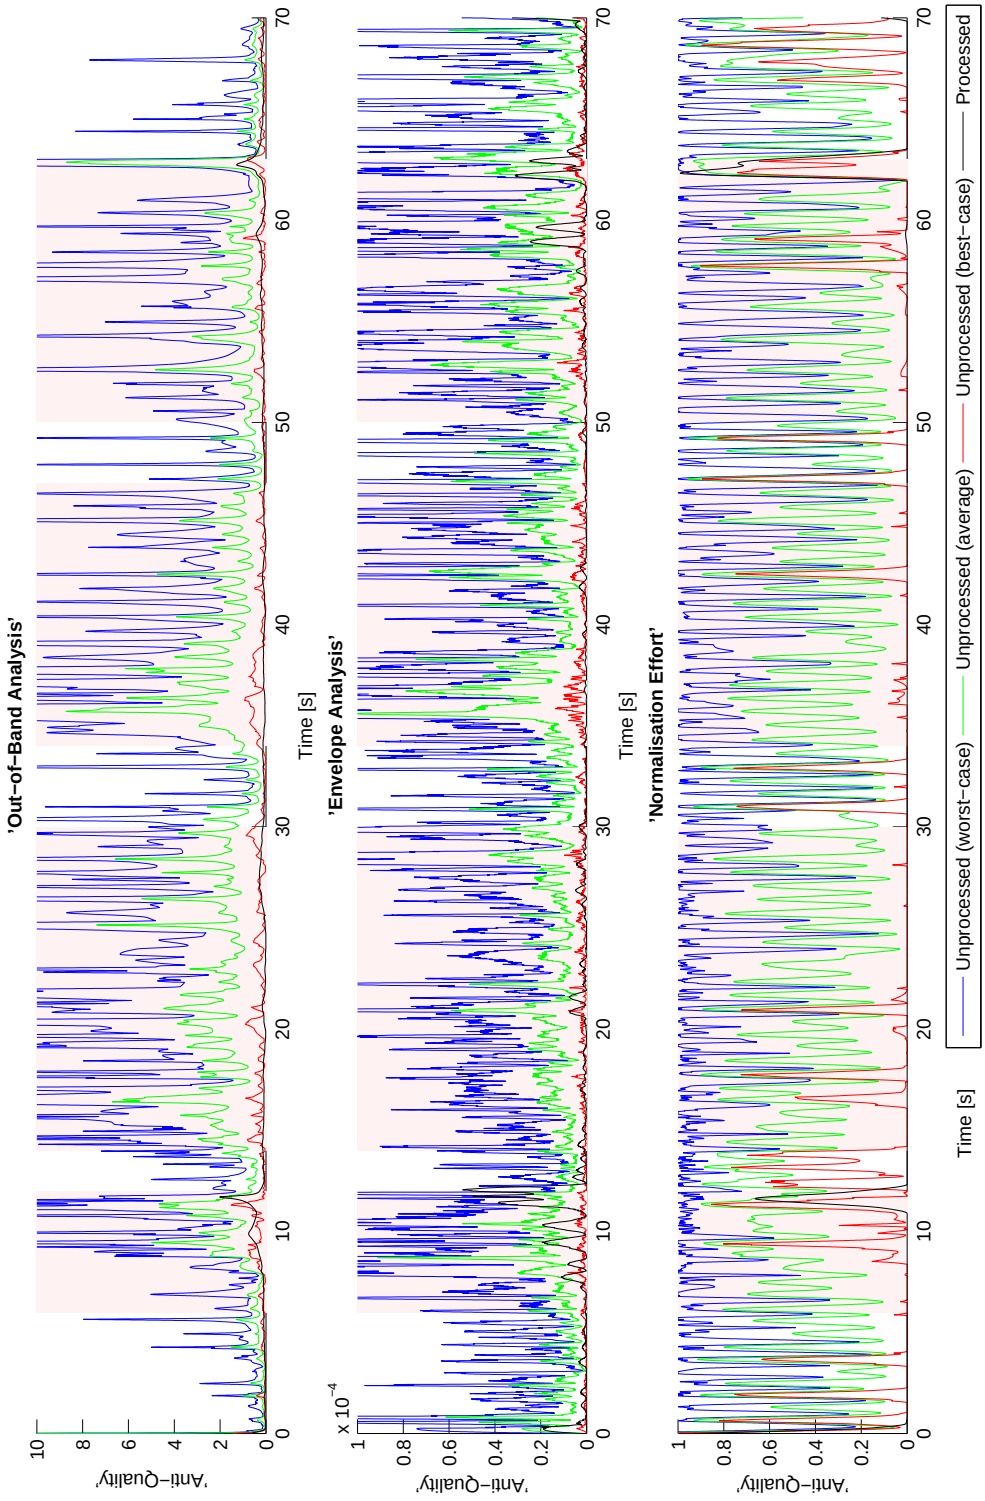

Figure 5.15: *'Qualities' of original and combined signals from participant 1's recording.*

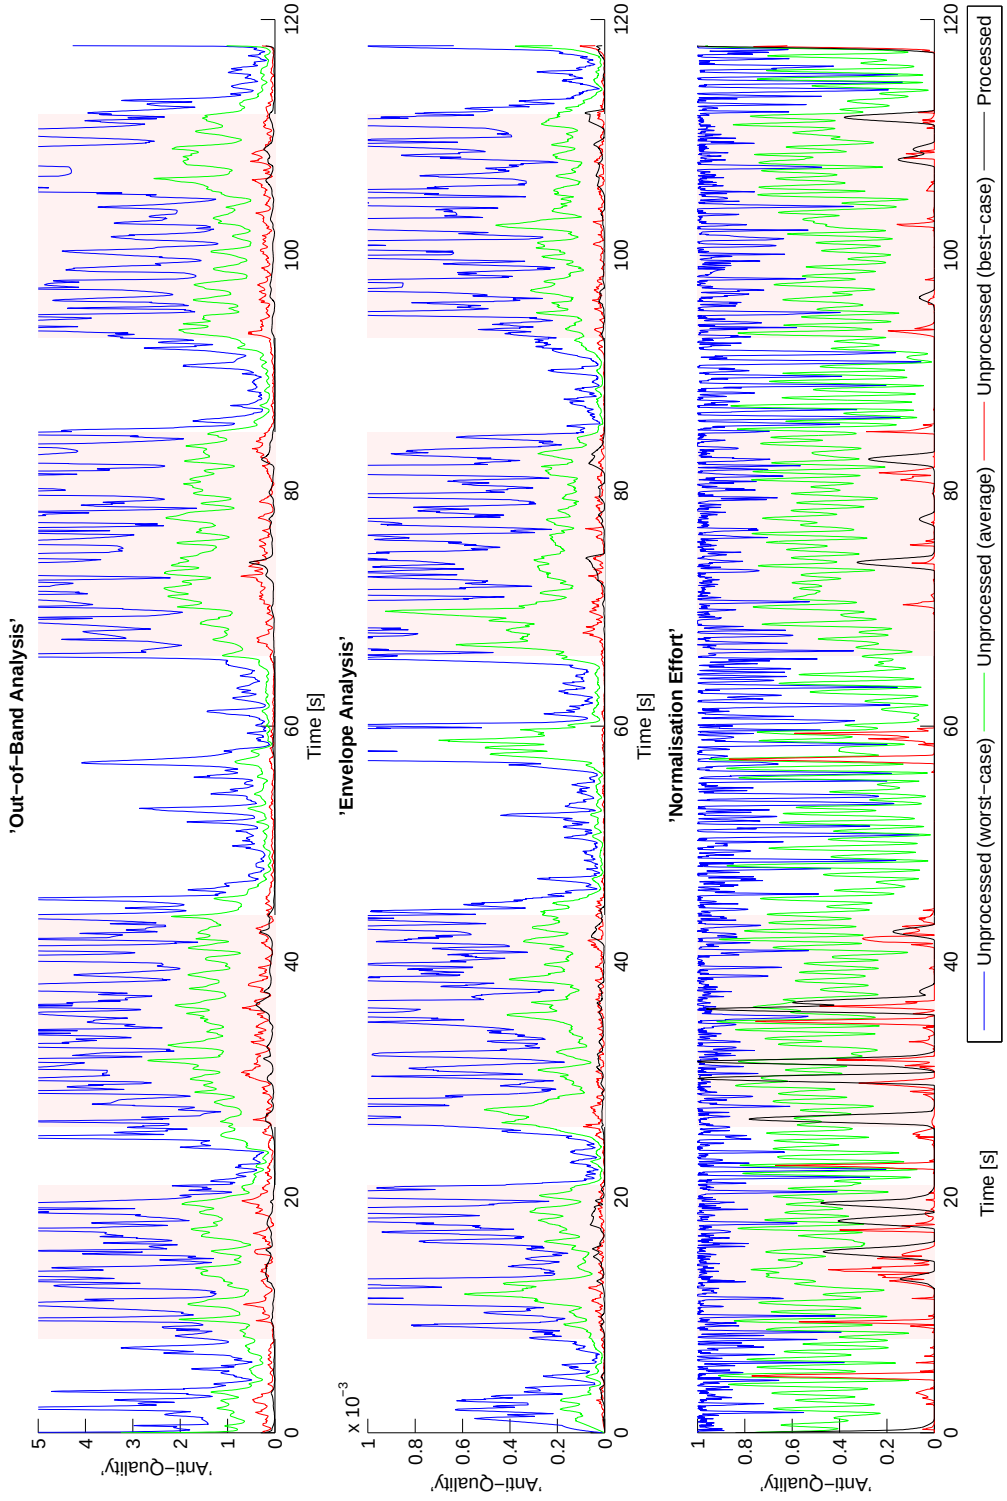

Figure 5.16: *'Qualities' of original and combined signals from participant 2's recording.*

5.4.3 Independent Component Analysis

As neither the pulsatile PPG nor the artefacts that occur during their measurement can be said to cause or directly 'affect' the other[5], statistically, they are said to be independent. This is one of the critical requirements of the signals when the technique called Independent Component Analysis (ICA) is used [93]. Additionally, the independent components' amplitudes must not have a Gaussian-like distribution. The PPG itself, which could be described as a sinusoid with a few additional features (the dicrotic notch, for example), does not have such a distribution. Repetitive and 'spiky' artefacts, too, have a non-Gaussian distribution. This makes ICA a potentially suitable method to use, in this case.

As ICA requires multiple inputs that are effectively linear mixes of the different components (of which one is the PPG), other researchers have previously tried using data obtained from a wide range of sources. Some have used an accelerometer as a direct, independent, measurement of the motion [94], with moderate success. However, as the motion artefact, as seen by the *optical* sensors, is not necessarily a good match with the *mechanical* movement, as seen by the accelerometer, there will be occasions where the algorithm may fail. One example of such an event is where motion occurs within the body (muscle contractions, etc), where the sensor itself is not displaced, but the optical signal changes.

Some researchers, in order to get around this limitation, have used purely optical techniques to obtain multiple signals; particularly using red and infra-red wavelengths that are common in pulse-oximeters [95] (additionally, researchers have used similar multiple-wavelength approaches using non-ICA techniques, for example [26]). The idea here is that PPGs (in reflection-mode) will have amplitudes that vary with different wave-

[5]In PPG acquisition systems, artefacts are often simply an additional signal that make it difficult (or even potentially impossible) to extract the PPG. The PPG still exists during their presence.

lengths due to the wavelength-dependent penetration depths [54]. Shorter wavelengths will not penetrate as deeply and will not contain as strong of a PPG, but the artefacts should remain. Generally, this has been shown to be an improved response over the accelerometer sources, but could still be improved as the different wavelengths will not 'see' all artefacts equally.

This is where the optical matrix has the potential to overcome these limitations. As all channels of the matrix use the same wavelength, then the signals that are observed *would* all 'see' both the artefacts and the PPG similarly between channels, if it were not for the mechanical coupling of, and the displacement between, the fibres. By mechanically coupling the fibres together, any motion that is imparted on one will also occur on others, however, as the physical positions of the fibres are different, the artefacts *may* not be perfectly matched: their amplitudes (not necessarily positive), will be unequal. The PPGs, however, are guaranteed to have the same polarity as they are generated by purely physiological (and not mechanical) sources. These linear mixes of PPG and artefact should make ICA an ideal method for separating them.

ICA is typically used when multiple components of the desired signal are *available* for separation, resulting in the desired signal being separated into several resulting component channels. For multi-probe ECG, for example, the well-known 'PQRST' complex is not just one 'component', but several that originate at different locations in the heart. Combining the resulting independent components will effectively reconstruct the 'pulse'. Whilst it is true that the PPG can be thought of as originating from two independent 'sources' (the two different sources of blood pressure as originally shown in Figure 1.2, described in Section 1.1.1.1), as the two pressures will effectively be summed when exiting the heart through the aorta, their effect will be 'seen' as a single component throughout the body. It is important to note that, in this case, the resulting channels that contain the PPG will all

contain the same signal. For example, if 16 different channels were available for extraction, there would only be one resulting component that contains the 'PPG' (although it may be represented in multiple channels).

Note: In the following example, the 'mixing matrix' has been ignored, and the calculated components are being used directly as the output signals. Normally, components that are seen to contain noise or other unwanted signals are 'zeroed' before being 're-mixed' to obtain the desired signal. The components are being used directly to illustrate the detected un-mixed components and how they are being considered (correctly) as independent by the algorithm.

All sixteen channels were fed into the `fastica` function (MATLAB), but with the output components limited to four due to the knowledge of the number of artefact types present. Of the resulting signals that were found (Figures 5.17 and 5.18), three show artefacts, and one shows a reasonably extracted PPG. It is interesting to note that of the three channels that show clear artefacts, they each show (roughly) a different 'type'.

For example, in participant 1's data, component 1 shows activity within region A, component 2 shows activity in regions B and C, and component 3 shows activity in region D. These three 'types' can be attributed to the three movement types. Similarly in participant 2's data, dominant activity is present in different channels, although less prominently.

The resulting component that primarily contains the PPG waveform[6] (IC 4), still shows a small artefacts at around 12 seconds for participant 1 and around 18 seconds for participant 2. These are both part of the compression movement that occurred, and may have similar properties to a PPG pulse to be distinguishable.

[6]For the purposes of this demonstration, the component that most dominantly contained the PPG signal (chosen visually) was moved to component 4. The ICA method that was used output the found components in a random order.

Although improvements to this technique are possible (additional pre- and post-processing, for example), the results at this stage are clear; by analysing the signals in order to extract/compute linearly mixed components (the methodology behind ICA), it is feasible to separate the PPG waveform from the underlying artefacts. There *may* be, however, situations where ICA will not work; particularly where the signal components, either PPG or artefact, have a near-perfect Gaussian distribution.

The 'compression' movement, surprisingly, was suppressed reasonably effectively, with it being extracted into component one of participant 1[7]. This would suggest that despite the supposed similarity of the artefacts between the different channels, there was a sufficient difference that allowed for the signals to be 'de-mixed' into their components. It is expected that some motions, where the observed variation is almost identical between all fibre channels, would result in it being impossible to separate the PPG and artefacts.

As the 'rotational' movements in the recording have shown that the 'centre of rotation' was not in the centre of the sensor (by the fact that complimentary pairs, as discussed, were found to be off to one side), the averaging method has not performed quite as well as the ICA. Despite the considerably more complex procedure of ICA, which would be difficult (but not impossible) to implement in a micro-controller, it does not require the amplitudes of the signals in order to extract the components. If the much simpler averaging were to be used, only complimentary pairs could be used without corrupting the resulting signal, whereas with the ICA method, all signals can be used, regardless of their amplitudes[8].

[7]Due to the `fastica` implementation in MATLAB using randomly selected initial conditions, the components will often change 'position' between successive runs (and also, between different participants' data).

[8]One of the pre-processing steps required for ICA to work is in-fact a filter that destroys (via normalisation) the original amplitudes and DC-components. This is the reason why the output components, likewise, contain no particular amplitude information (relative to the inputs), and can even appear inverted

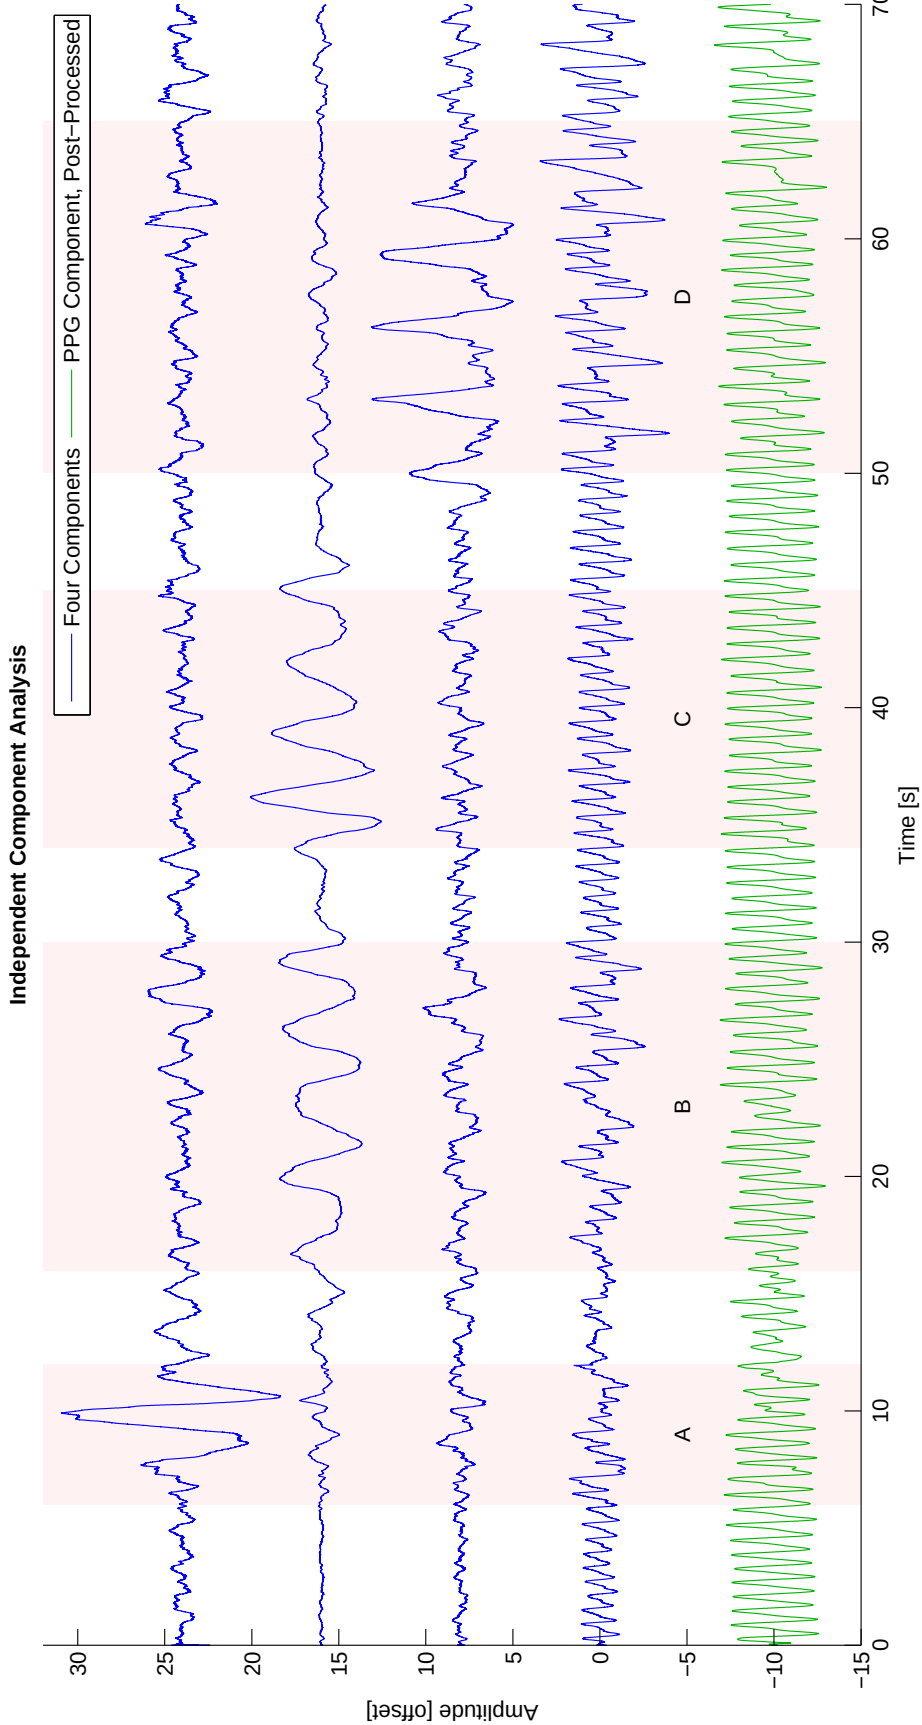

Figure 5.17: *Participant 1: Four extracted independent components, three of which show artefacts (IC 1-3), and one that primarily contains the PPG (IC 4). A normalised and filtered (Butterworth, 0.66-9.00 Hz) version of IC 4 is also shown (bottom). Four red bands illustrating the regions of the artefacts is also present.*

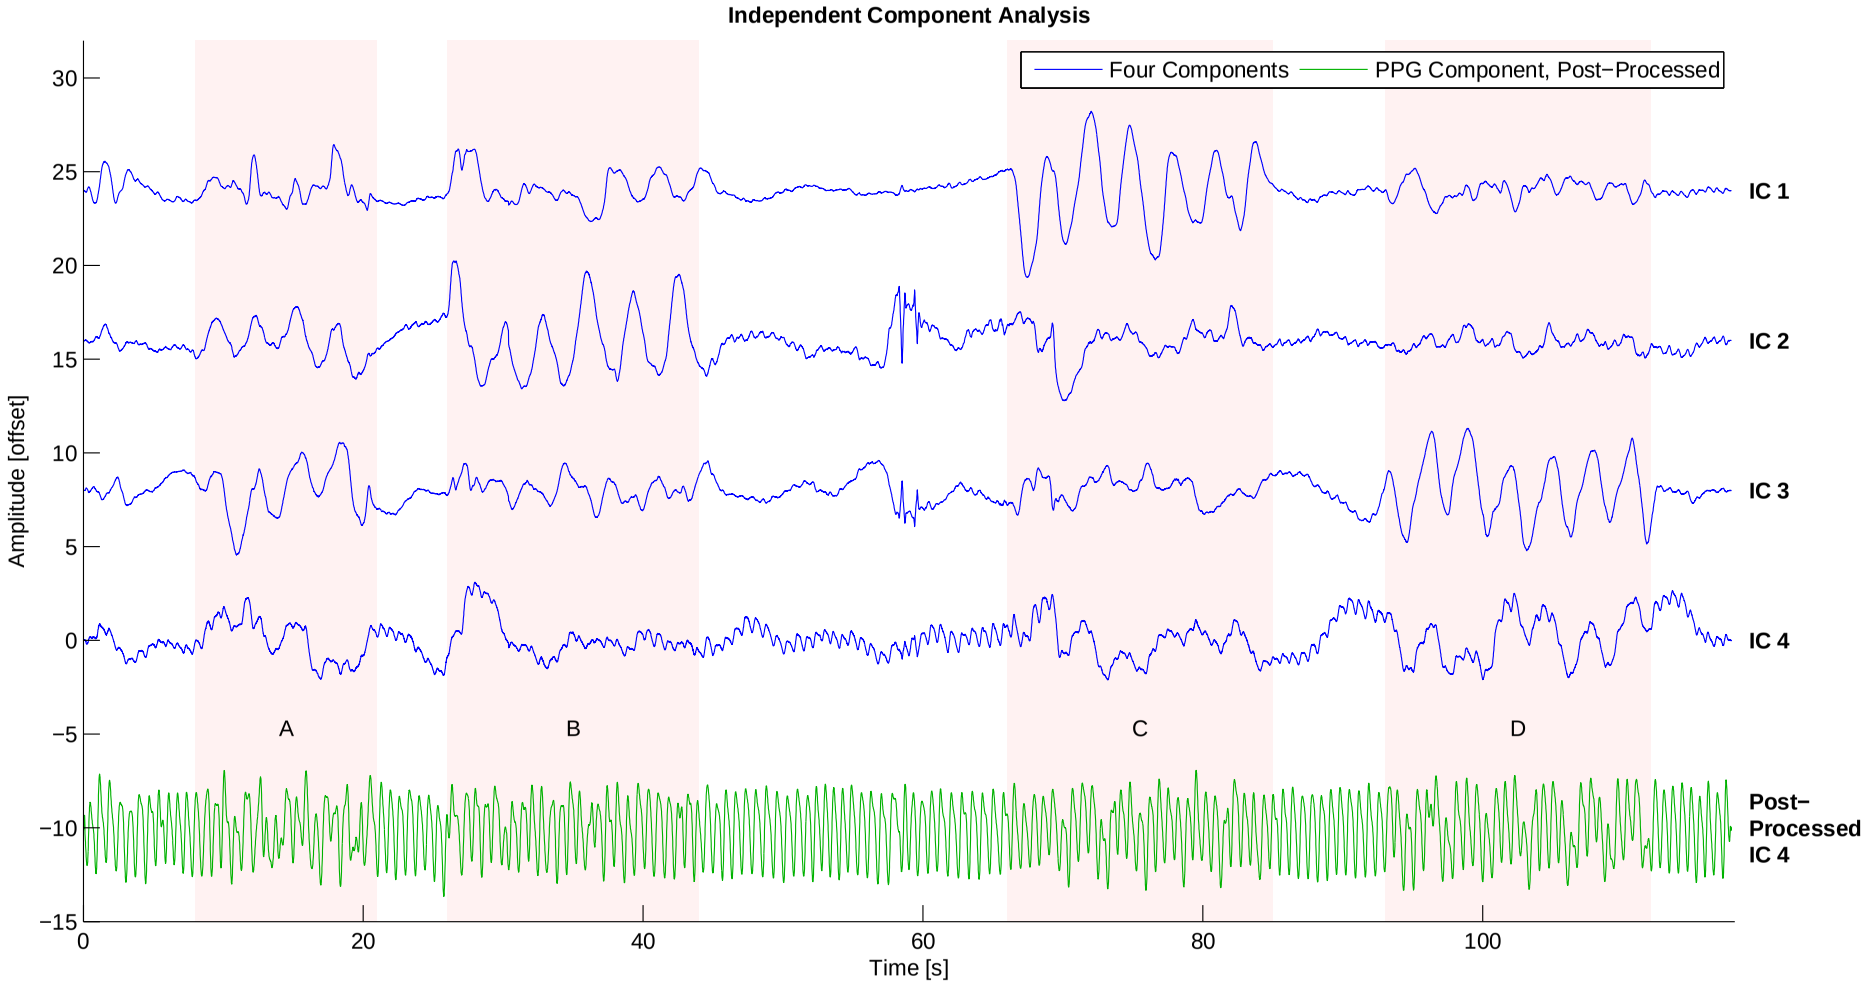

Figure 5.18: *Participant 2: Four extracted independent components, three of which show artefacts (IC 1-3), and one that primarily contains the PPG (IC 4). A normalised and filtered (Butterworth, 0.66-9.00 Hz) version of IC 4 is also shown (bottom). Four red bands illustrating the regions of the artefacts is also present.*

5.4.3.1 'Quality' Algorithm Results

In an identical method to the averaging algorithm, by passing the two participants' raw data through the previously developed 'quality calculating' algorithms, and comparing the results with the results of the processed data (average of the sixteen channels in this case), the effectiveness of the processing technique can be quantised.

Figures 5.19 and 5.20 show the results of the raw (16 channels) and the processed (ICA) data after the three 'indirect' quality measuring functions: 'out-of-bounds', 'envelope', and 'normalisation effort' analysis.

Again, to simplify the graphs, of the sixteen channels the algorithms were run on, the *worst*, *average*, and *best* 'qualities' are shown. These are presented in blue, green, and red, respectively. The result from the processed data is shown in black.

In both participants' recordings, the regions where no movement was present shows a significant 'drop' in the 'anti-qualities' - more-so than in the channel-averaging results. The algorithms represented in the upper two graphs show a larger number for more 'chaotic' (lower quality) signals. The algorithm for the lower graph normally represents a true 'quality' (from 0 to 1), but has been inverted so as to keep consistent with the other graphs.

Statistically, none of the quality vectors from any of the three algorithms are normally distributed, as tested using the One-sample Kolmogorov-Smirnov test (with $p < 0.05$) [91]. This means that a valid comparison between the 'best channel' quality and the processed 'averaged' channels could be made using a Wilcoxon signed rank test [92]. For the 'out-of-bounds' and 'envelope' algorithms, the chosen ICA component of the sixteen channels yields a statistically different quality result ($p < 0.05$) than the best-case channel.

Tables 5.4 and 5.5 show the times-improvements of the chosen ICA component of the channels, over the best-channel quality results, for the

'out-of-bounds' and 'envelope' analysis algorithms, respectively. Just as before, the 'normalisation' analysis algorithm cannot be included in these statistics tests as more than 50% (for each participant) of the resulting quality data is identical, and thus cannot be declared as statistically different using the described method. Table 5.6 shows the percentage of time each algorithm outperforms the other.

Over the two participants, understanding that a quality level of 0 represents a perfect PPG, with a positive value representing a degraded signal, the 'ICA algorithm' produces - on average between the participants - a 30.4× improvement in measured quality (using a combination of the out-of-bounds and envelope methods).

The average median of the participants data for the 'ICA algorithm' is more than ten times lower than that of the 'averaging' algorithm.

Participant	**ICA** *Median (IQR)*	**Best Channel** *Median (IQR)*	**Improvement** *(best/ICA)*
1	0.280 (0.237)	1.321 (2.135)	4.710 ×
2	0.094 (0.127)	1.363 (1.715)	14.471 ×

Table 5.4: *'Out-of-bounds' algorithm output statistics, compared to the 'best-case' channel, for both participants. All medians/IQRs* $\times 10^{-1}$.

Participant	**ICA** *Median (IQR)*	**Best Channel** *Median (IQR)*	**Improvement** *(best/ICA)*
1	0.0044 (0.0268)	0.098 (0.158)	22.369 ×
2	0.0122 (0.0361)	0.978 (1.422)	79.958 ×

Table 5.5: *'Envelope' algorithm output statistics, compared to the 'best-case' channel, for both participants. All medians/IQRs* $\times 10^{-5}$.

Participant	**I >B** *(%)*	**I = B** *(%)*	**I <B** *(%)*
1	39.6	52.7	7.7
2	38.2	57.7	4.1

Table 5.6: *'Normalisation' algorithm output statistics, compared to the 'best-case' channel, for both participants.* ***I****: ICA algorithm 'quality',* ***B****: Best channel quality. Where higher 'quality' is better (for example, 'I >B' represents the ICA algorithm outperforming the best-channel result).*

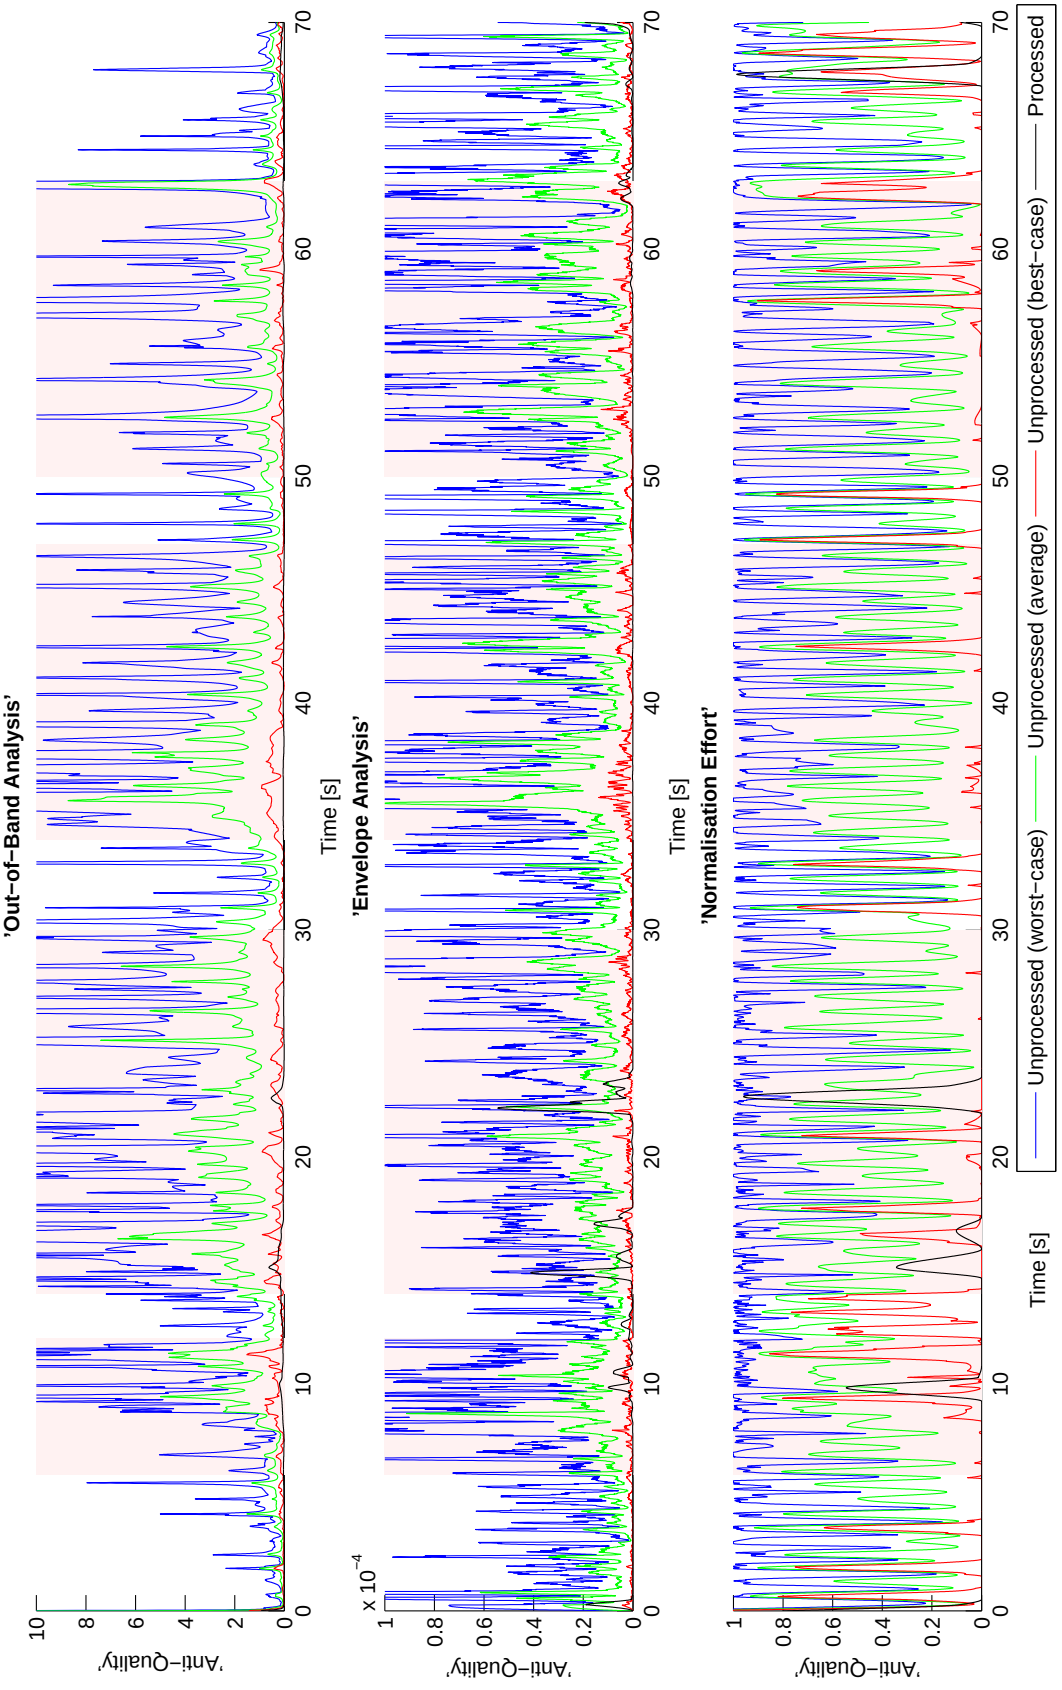

Figure 5.19: *'Qualities' of original and combined signals from participant 1's recording.*

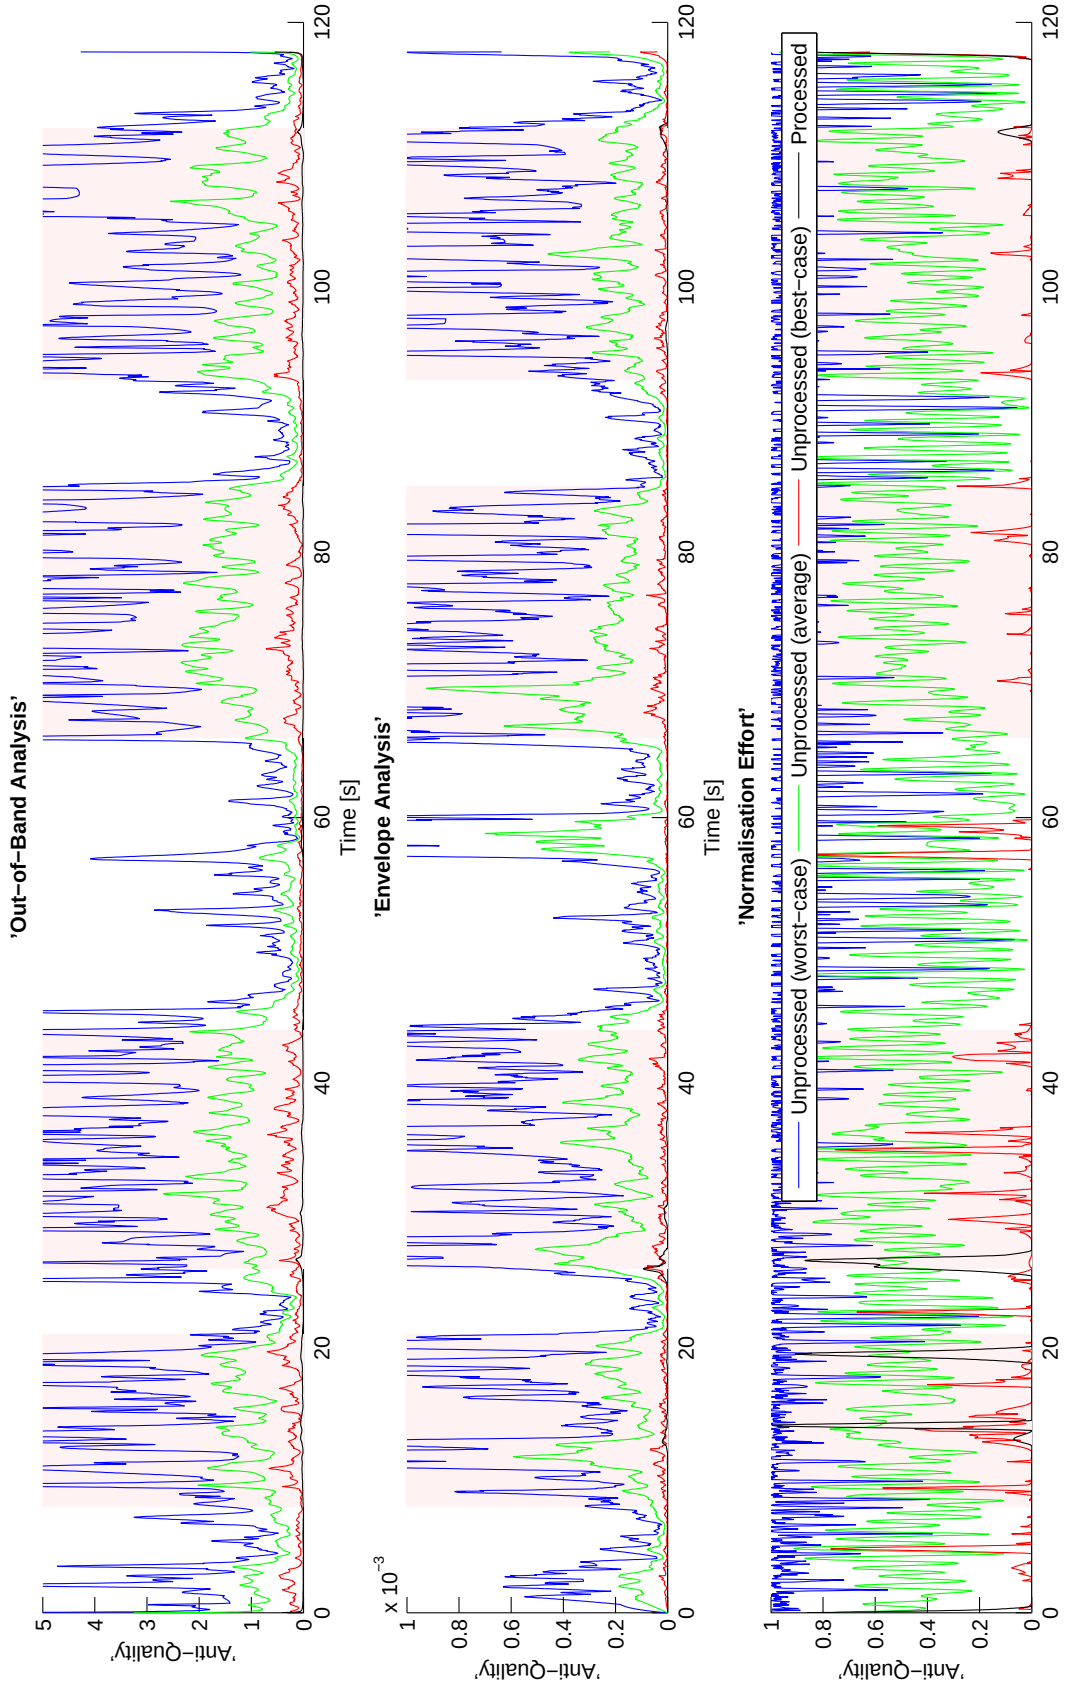

Figure 5.20: *'Qualities' of original and combined signals from participant 2's recording.*

5.5 Discussion

It has been shown that, for the 16-channel 'optical matrix', improvements to the resulting signal quality can be made by utilising redundant information provided by the different channels. If a simple 'averaging' of the channels' waveforms is performed, then a typical (average) 2.5× improvement in measured quality can be achieved (using a combination of the out-of-bounds and envelope methods). For the independent component analysis technique, a considerably better 30.4× improvement in measured quality is possible (using a combination of the same methods).

The design of the 'Optical Matrix', which would not have been conceived without information learned from the previous non-contact work, has shown to be highly successful, even for the limited datasets that were obtained. Having multiple channels that are mechanically coupled together, but span a larger measurement area than just a single sensor, allow for artefacts that manifest in the signal to appear with different amplitudes and polarities over the sensing region. These differences are enough to separate the PPG from the artefacts using either simple averaging when the artefacts in neighbouring channels have differing polarities, or by using independent component analysis (a more computationally taxing method) for when amplitude and polarity information does not vary as much.

One limiting factor of the presented techniques, however, is when the artefact presents itself as a 'compression' perpendicular to the surface of the skin. In this 'mode', all channels are subject to a similar artefact (amplitude and polarity), and so cannot be easily separated from the PPG signal, particularly when the artefact and signal are similar.

Chapter 6

Summary, Conclusions and Future Work

The research described in this thesis has revealed a series of interesting and important pieces of knowledge that have been critical in achieving the final goal of reducing the effect that motion has on the PPG signal. In addition to work that could potentially be carried out in the future, including some progress that has been made by another researcher based on the results of this project, key findings and outcomes are summarised next, before concluding with how, and to what extent, the main goal has been reached.

6.1 Summary and Conclusions

One important idea that was developed during the research, that was critical in the development of the final device, was the idea that any means of measuring the motion itself, that did not use the same method as the PPG extraction, would ultimately *not* necessarily be sufficient to 'correct' for the error, or at the very least, provide a poor result. This was simply due to the difference in magnitude between the small PPG component, and the comparably massive artefacts; and was observed from the results of the experiment with the optical matrix, whereby the artefacts within the signals from neighbouring fibres were significantly different *despite them being detected from mechanically-coupled channels with identical specification.* As the small distance between the adjacent sensors was sufficient to create such

differing results, it would be reasonable to think that other methods of detecting the motion that have *less* of an association with the PPG (such as using accelerometers, etc.) would result in even more significantly differing results. Some specific examples of how prior researchers have used such 'alternative' methods to attempt to selectively filter or recreate the artefacts were discussed in Section 2.4. By using multiple channels that are optically and electrically identical in order to make a single measurement, with each source/sensor pair outputting a signal which contains both the PPG and artefacts, several 'perspectives' of the same region can be observed. Whilst the PPG signal is reasonably consistent between all channels, any artefact will be seen differently depending on the 'direction' of the motion. The very fact that both components will be present, to some extent, in every channel means that calculations on the signals may be able to yield the components via separation (as the independent component analysis has shown).

An additional processing step that was determined to be necessary was the technique required to correctly filter PPG signals, such that there is minimal loss of *information* whilst removing unnecessary *noise*. Whilst a narrow band-pass filter could adequately remove all but the PPG (with or without its harmonics), this may not be particularly useful as there is additional information present in the outlying signal that can be used. Section 2.4.4.1 discusses this concept, and provides a tabulated list of other researcher's filters (that are often too narrow to be useful for anything other than displaying the data). As the motion artefact is, by the definitions used in this research, a signal, instead of noise, then if it were to become the desired output, then the PPG would in turn become the erroneous component. This interchangeability of focus of the two components, and the fact that the two signals have often very similar frequency ranges, means that the artefact must be considered as an independent signal in the resulting data: one that must be *separated* out, instead of 'removed' via

filtering. If multiple signals are available (as is the case with the optical matrix), then by using certain characteristics of the PPG and artefacts that would not have been visible if overly-narrow band-pass filters were used, then the presence of each component in the signals can be calculated and output (by the means of a 'quality' measure). If a processing algorithm requires it, particularly with heart-rate extraction techniques, the resulting quality index can be used to select channels or signals that have sufficiently strong PPG components (or weak artefact components) in order to improve the probability of it working successfully. Ultimately, this quality-level calculation would not be possible with overly filtered signals; indeed, even white noise that has been filtered with the incorrect filter can (and will) appear to look like a smooth PPG signal, which has been demonstrated.

Unfortunately, perhaps, the results of the research presented have concluded that a single sensor design may not have sufficient redundancy in the acquired data to accurately extract a PPG when artefacts are present in the signal: multiple sensors are required. Ultimately, this means that existing devices cannot be 'upgraded' with new sensor heads and firmware, for example, as the electronics that drive them are unlikely to be able to handle so many channels. Pulse oximeters typically have two wavelengths that are used to calculate the blood-oxygen saturation; whilst a lower bound of how many channels are required to separate artefacts from the PPG has not been explored in this research, it is predicted that at least three are required (giving two orthogonal pairs, with one common channel), although more are predicted to be required. On the positive-side of this, a new product may be viable to exploit this technique (or at least the technique may be used as part of new designs), which may make this development tempting from a commercial point-of-view. More on this is in the 'Future Work' section that follows.

The optical matrix, as presented, has been demonstrated to work well

with a single wavelength (green, ≈550 nm), making it suitable, in its current state, for just photoplethysmography. One primary advantage of this technique is that, fundamentally, each channel behaves just like a pre-existing single-element sensor. In other words, there is nothing specific about the design that would prevent there being multiple wavelengths used. In this case, each LED in the array can be replaced by two (or more) LEDs with different wavelengths, each individually behaving in the same way. By multiplexing between the wavelengths using, say, time-division-multiplexing (TDM), the application of photoplethysmography can be extended to pulse oximetry. Although reflection-mode pulse-oximetry is not used as widely, there are advantages to it, especially when used on an ill patient who may have poor blood circulation at the peripheries and so transmission-mode variants will have poorer performance [96].

A vast majority of clinical photoplethysmography systems are based around a large 'base-station' with a wired connection to a small sensor (a finger-clip, or disposable patch), making the analogue electronics to manipulate the signal from the sensor, and the digital electronics to process the results, not have to be made either small or low-power. For these particular applications, there may not be any advantage to 'miniaturising' the optical matrix (other than the fibre bundle) as almost the entire acquisition and processing back-end could be located in a similar 'base-station' on a rack in a hospital. In this case, the fibres themselves could simply be made smaller in order to make the bundle more flexible, with a better interface established with the optical components in order to minimise losses. However, there is an increasingly common usage of photoplethysmography, particularly in reflection-mode, that requires both the sensor head and the processing unit to be not only considerably smaller, but low power: wearable technology. For health monitoring during fitness and sport activities, through to medical home-care, the PPG system (that most often takes

the form of a 'smart watch' or wrist-band) has to be small. For these applications, where the heart-rate is monitored during physical activity, the resilience of the system to motion must be greater. It is here where the developed system would need to be made considerably smaller. Possible avenues of miniaturisation are presented in the 'Future Work' section that follows.

During the development of the processing algorithms for the optical matrix, the finding that independent component analysis of the individual channels was able to separate out the 'compression' type of motion was unexpected, as it was assumed that a greater variation between the source channels was required. The similarity of the artefact component in each of the channels was revealed when the channels were averaged, and the compression artefact was not removed. This indicated that more sophisticated processing algorithms (such as ICA) are required for certain types of artefact, whereas other types can be suppressed simply by using their differential nature, and be 'averaged out'. Whilst there is a small possibility that there exists a small differential component spread between the channels that enables the ICA to separate the artefact from the PPG, as the same effect was seen in all the participants' recordings, this would suggest that the dissimilarity between the PPG and the artefact was sufficient for the algorithm to work. The largest issue when using independent component analysis is converting the 'batch-processing' (post-processing) style of algorithm into a real-time one suitable for obtaining the results *during* the acquisition, as would be required in most applications. This is certainly possible, as it has been applied to other types of data previously [97], however, this has not been explored in this research. The combination of the added computational complexity required to separate the PPG from other signals, and the extra light sources used in the developed hardware, will ultimately increase the power consumption of the system. For portable

devices, this means that either the total operating time (battery life) is decreased, the batteries that are used must have increased capacity, or the protocol by which the LEDs are driven is changed. A few possibilities are presented later.

6.1.1 Concluding Remarks

The primary aim of this research, as the title states, was to create a method that reduces the negative effects that motion has on reflection-mode photoplethysmography. Whilst other researchers have focussed on countering motion artefacts by using novel algorithms on existing signals, bespoke mechanical designs to improve the optical coupling between the skin and the sensor, and attempting to 'subtract' the artefacts by measuring the motion by other means, all research for in-contact methods have focussed on using just a single light sensor. The journey that this research took explored not only these ideas, but other possibilities, including a design that has not been presented before in the literature. A large amount of inspiration came from observing the data that was recorded from non-contact imaging sensors (cameras), whereby the signals from different regions of the same subject appeared different, yet contained a similar PPG signal.

The end-result of this exploration has been an early prototype of a novel sensing array, similar in many ways to an imaging camera, but using the in-contact methodology, and with combined light sources interleaved with the sensors. By using redundant information in a larger set of data, the PPG signal and any artefacts that are present can be separated using a standard algorithm (ICA). The resulting components can be selected using developed techniques that estimate the 'quality' of the signal, and the heart-rate can be obtained. The result is a system that has ultimately been demonstrated to reduce the effect that motion has on the obtained PPG signal.

6.2 Future Work

As the intention of this research project was *not* to create a marketable product, more-so to simply explore techniques to reduce the effect of motion, the resulting design of the optical matrix is physically large, delicate, not portable, impractical and power-hungry. In essence, it is only suitable as a prototype to demonstrate the effectiveness of the technique. This section presents ideas (and some progress) to reconstructing the system into a more 'user-friendly' and useful device.

6.2.1 Integrated, Simplified System

Currently, in order to use the optical matrix, a large sequence of manual steps must be taken. The Raspberry Pi must first boot, and the software on the PC must establish a connection to it. A recording can then be made which is stored in the PC's main memory (RAM). After manually saving the data to a file, it can later be processed in MATLAB (using the ICA algorithm, for example) to extract the independent components, each of which are then run through the individual quality-level algorithms in order to predict which of the components is a PPG; the 'by-product' of these algorithms is the heart-rate and 'quality'. The manual organisation and control of the data and equipment makes this overly complicated, and can be greatly simplified.

Ideally, all processing should be performed on the device/system itself, and should not rely on an external computer, as this would not make it truly portable; this would necessitate the inclusion of some kind of processor. A new system, if designed correctly, could remove the need for the FPGA and Raspberry Pi, and instead simply use a single processor that generates the necessary patterns for the LEDs, acquires the resulting signals via the ADCs, performs the processing (real-time ICA, and the quality-level

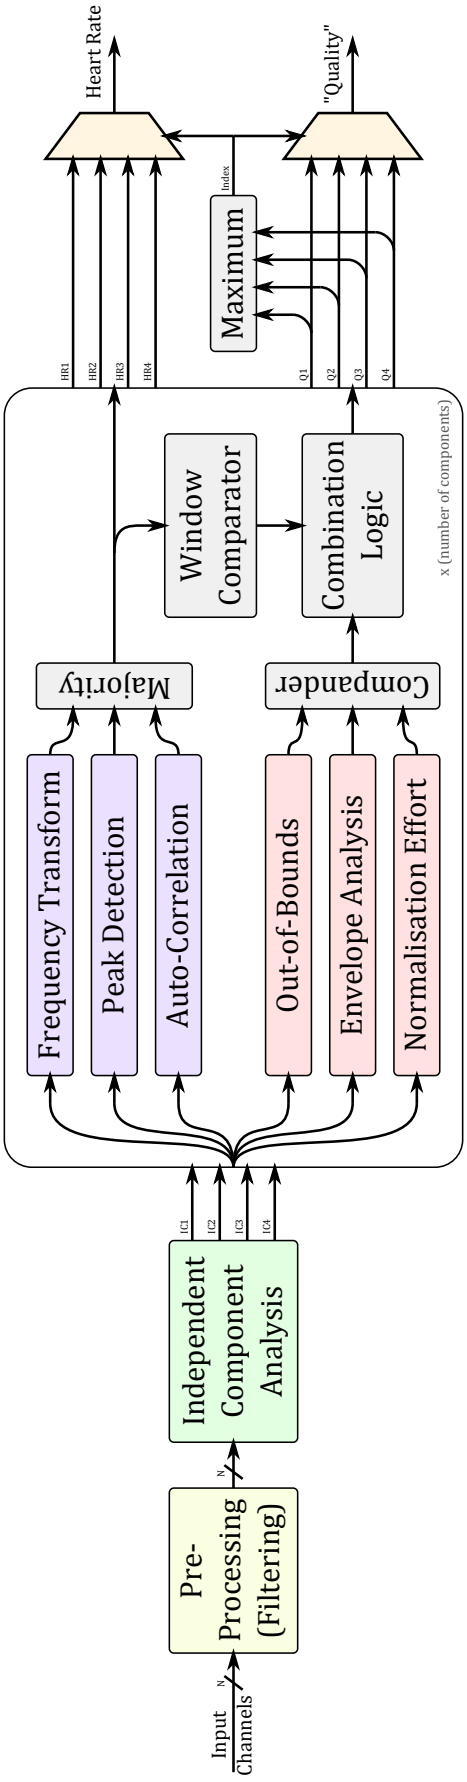

Figure 6.1: *Illustration of the processing steps required to extract a PPG (and other information) using the optical matrix's sixteen channels.*

measurements), before outputting just the PPG component and heart-rate. The processing steps that are required are illustrated in Figure 6.1. As the 'user-friendly' device does not need to be 'developer-friendly', the raw data is not required to be output, so high-bandwidth communication links are not necessary[1]. Developing this sparse system into a single integrated one would be an interesting project, just in itself.

By reducing the number of hardware components, and combining the required data processing into a single 'space', the overall size of the system can be dramatically reduced, and the hardware made robust. This miniaturisation and simplification also allows the device to potentially be portable, as the data that could transmitted from it is minimised (a single low sample-rate PPG waveform, and heart-rate data, instead of sixteen high sample-rate signals). This would also minimise the power consumption, making battery power feasible.

6.2.1.1 Low Power Operation

In one mode of operation of the optical matrix, all sixteen LEDs are illuminated to a constant brightness; it was this mode that gave the results presented in Chapter 5. The internal sampling time of the ADCs is considerably smaller than the duration of one sample period without degrading the signal, in which case, the LEDs need not be illuminated for that time. In fact, only a small lead-in time, where the LED is illuminated before the sampling starts, is required to allow for the system to reach steady-state (ringing to dampen, for example). Even if the ADC sample times were around 1 ms (this is a **very** long time for sample-and-hold circuitry), for a 100 sample-per-second system, the LEDs could be extinguished for almost 90 % of the time. In reality, an ADC sample time could (and would) be must less than this, giving a much reduced power consumption.

[1] Although, if a sufficient bandwidth were available, the raw and intermediately processed data could be output, for further development or research purposes.

Appendices

I Envelope Algorithm

In order to quickly and reliably calculate the envelopes of signals, a simple function was written that performed the task using two stages. The first stage behaved much like a lossy peak-detector, and the second, an interpolator.

For the lossy peak detector, each sample of the signal is compared with a variable (say, x) which has been previously initialised to the first value in the signal. If the incoming sample is greater than x, then x is set to this new value. In other words, for a monotonically increasing signal, x *equals* the signal. Now, if the next sample is less than x, then x is multiplied by a 'decay constant', which is a value less than 1. This way, if the signal 'falls away' from x slowly, x will follow the signal. However, if the signal falls quickly, the x will remain behind. The value of x represents the real-time 'lossy peak' of the signal (*similar* to a channel amplitude visualisation in a graphic equaliser).

In order to make the first-stage envelope symmetrical, the algorithm is also run on the reversed data (this requires some amount of buffering of the data, if the incoming signal were actually streamed to the algorithm). The maximum of the two detected envelopes is then calculated, which results in the first stage signal (the upper red line in Figure 2). In order to obtain the first-stage envelope that follows the *troughs* of the signal, the user of the function should input the inverted signal, and then invert the result (this yields the lower red line in Figure 2).

The second stage envelope, which is calculated within the function, is an

interpolation of the peaks of the first-stage envelope. Because of the nature of the first-stage envelope, small fluctuations in the signal that lie on the rising and falling edges are not picked up, fluctuations that would otherwise likely be picked up as additional peaks. Because these 'peaks' do not exist in the first stage envelope, the interpolated signal faithfully follows the true peaks in the signal, allowing for a smoother envelope to be formed. It is this new interpolated signal which is output as the second-stage envelope.

The MATLAB function that was used is presented in Listing 1.

```
function [env, int] = findenv(x, decay, interp)
%FINDENV Find envelope of data
%   [ENV,INT] = FINDPEAKS(X,DECAY) finds two envelopes of the
%   signal (X) using the decay constant (DECAY). The first
%   envelope (ENV) represents the raw decay-style envelope.
%   The second, (INT), represents the interpolated signal
%   that is more akin to the analytical envelope.

%   Matthew Butler 2013-2016.

% Default to pchip interpolation if none given.
if nargin < 3
    interp = 'pchip';
end;

% Allocate space for first-stage envelopes.
env_fwd = zeros(length(x), 1);
env_rev = zeros(length(x), 1);

% Initialise the tracking value.
curr_fwd = x(1);
curr_rev = x(end);

% Iterate over each element.
for i = 1 : length(x)
    % Forward: update or decay.
    if x(i) > curr_fwd
        curr_fwd = x(i);                % More (update).
    else
        curr_fwd = curr_fwd * decay;    % Less (decay).
    end;
    % Reverse: update or decay.
    if x(end-i+1) > curr_rev
        curr_rev = x(end-i+1);          % More (update).
```

```
    else
        curr_rev = curr_rev * decay;    % Less (decay).
    end;
    % Save current values.
    env_fwd(i) = curr_fwd;
    env_rev(end-i+1) = curr_rev;
end;
% Maximum of the two envelopes.
max_env = max(env_fwd, env_rev);

% Use built-in peak-finder.
[peak_value, peak_index] = findpeaks(max_env);

% Pad ends to keep interpolation happy.
peak_index = [0; peak_index; length(x)+1];
peak_value = [0; peak_value; 0];

% Output first-stage envelope.
env = max_env;

% Interpolate first-stage peaks to get second-stage.
int = interp1(...
    peak_index, ...     % Input X.
    peak_value, ...     % Input Y.
    1 : length(x), ...  % Output X.
    interp, ...         % User defines interp method.
    'extrap');          % Extrapolate if necessary.

end
```

Listing 1: *Developed 'decay-style' envelope detection algorithm.*

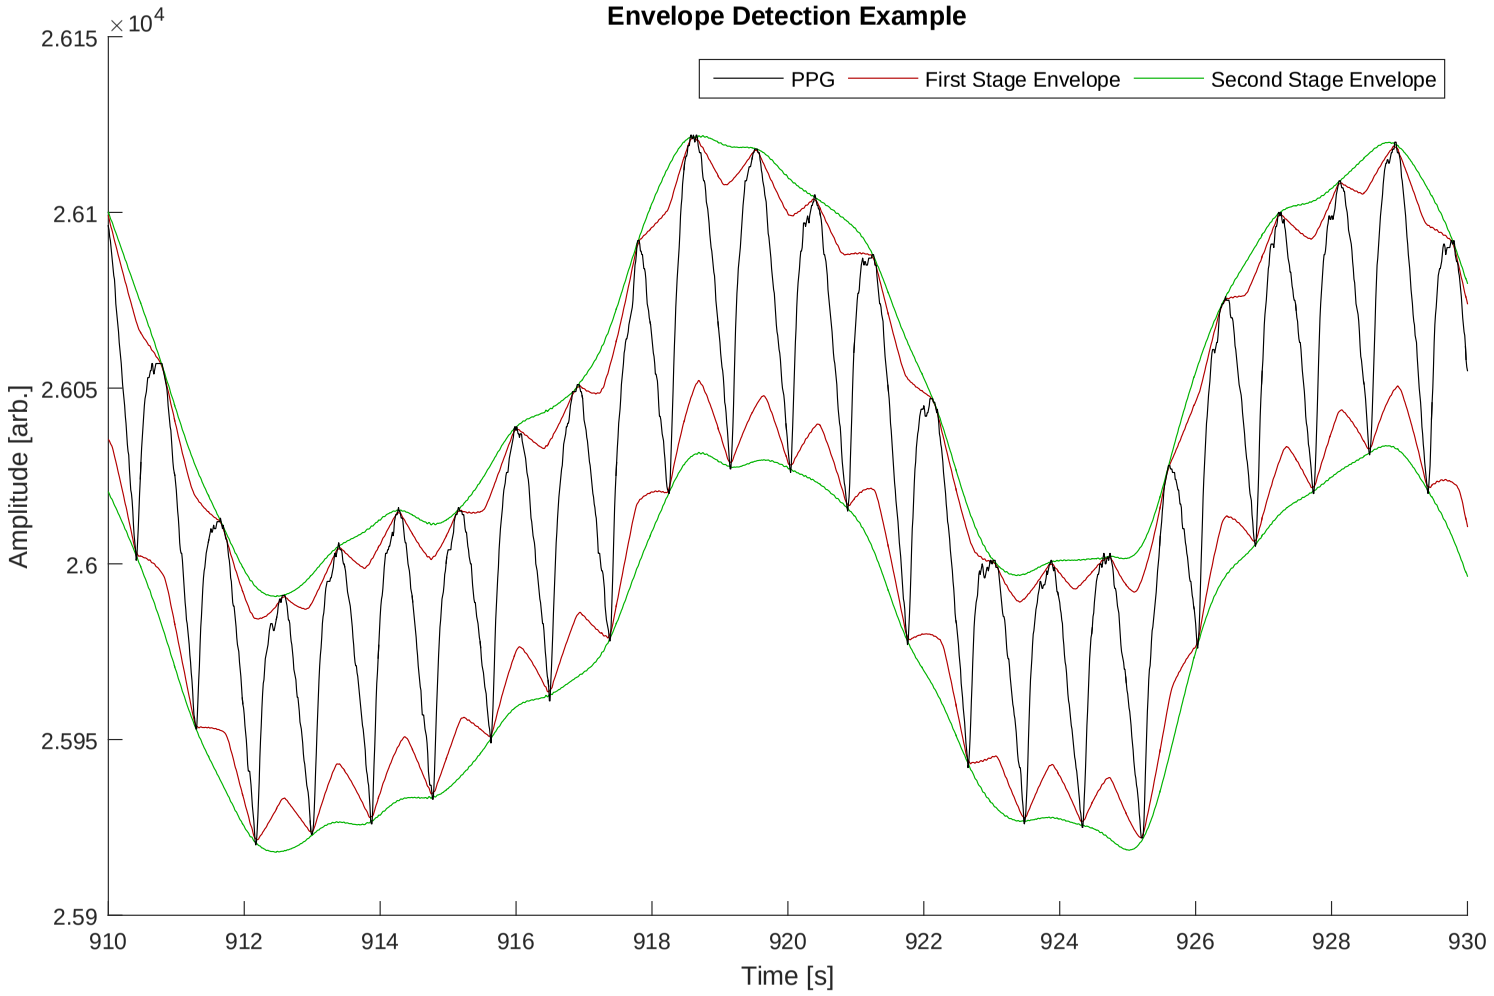

Figure 2: *Illustration of the two output stages of the envelope detector that was developed for some of the algorithms in this research. The 'second-stage' envelope is an interpolation of the detected peaks of the first-stage.*

II PPG-ECG Recorder Schematic & PCB

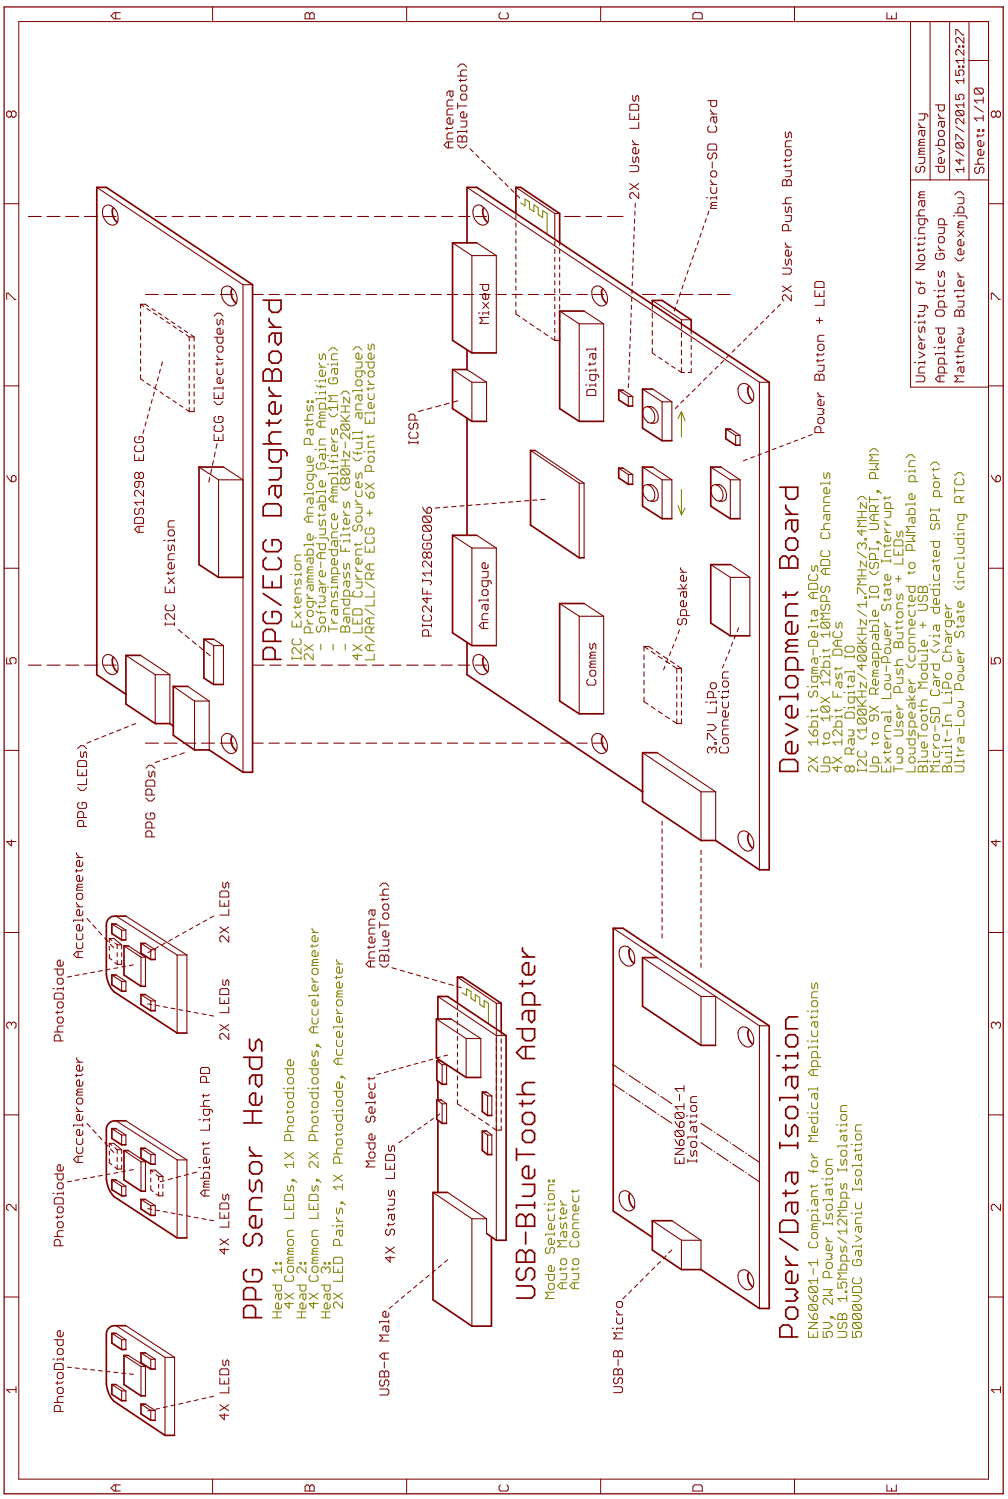

Figure 3: *Design overview.*

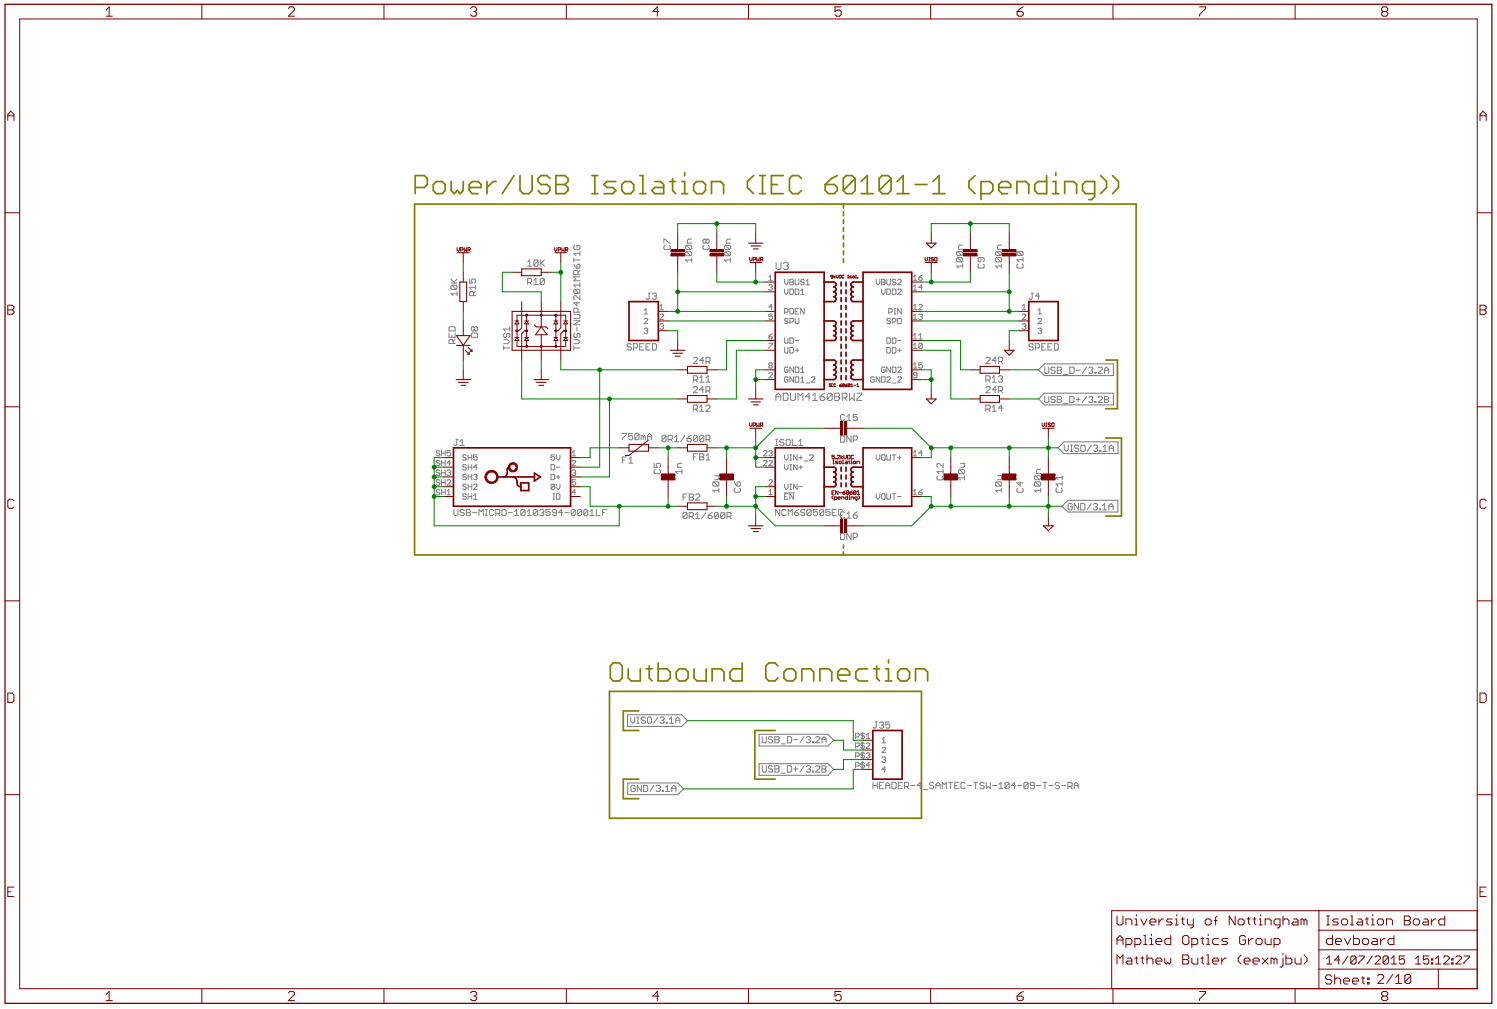

Figure 4: *Isolation Board (power/data coupling).*

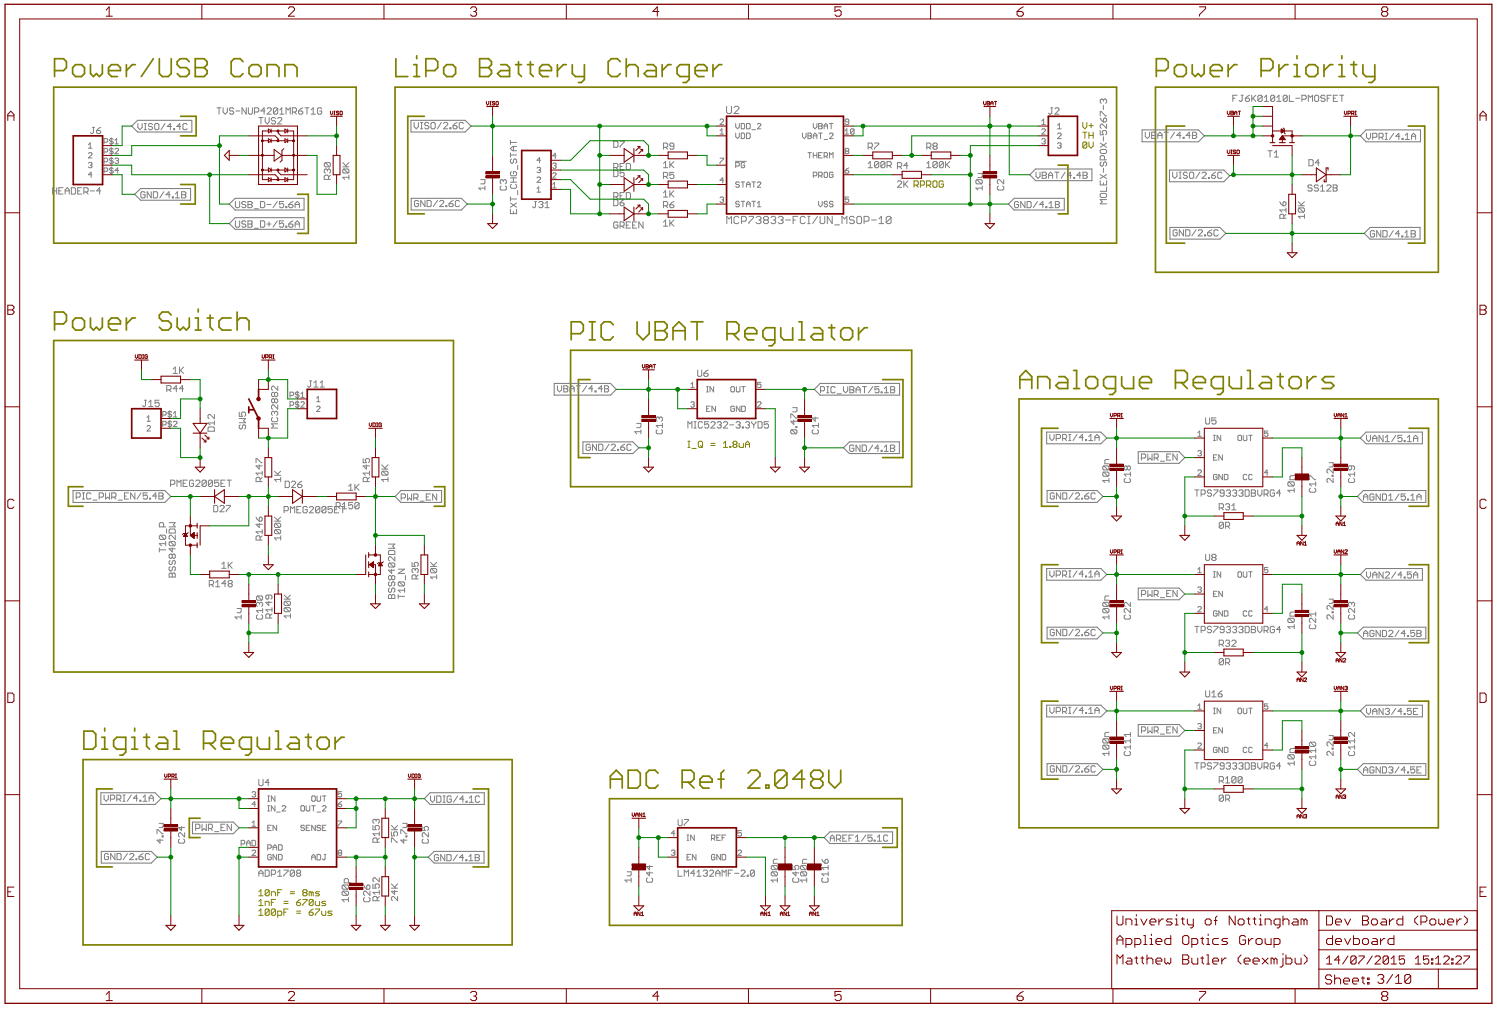

Figure 5: *Power regulation and control.*

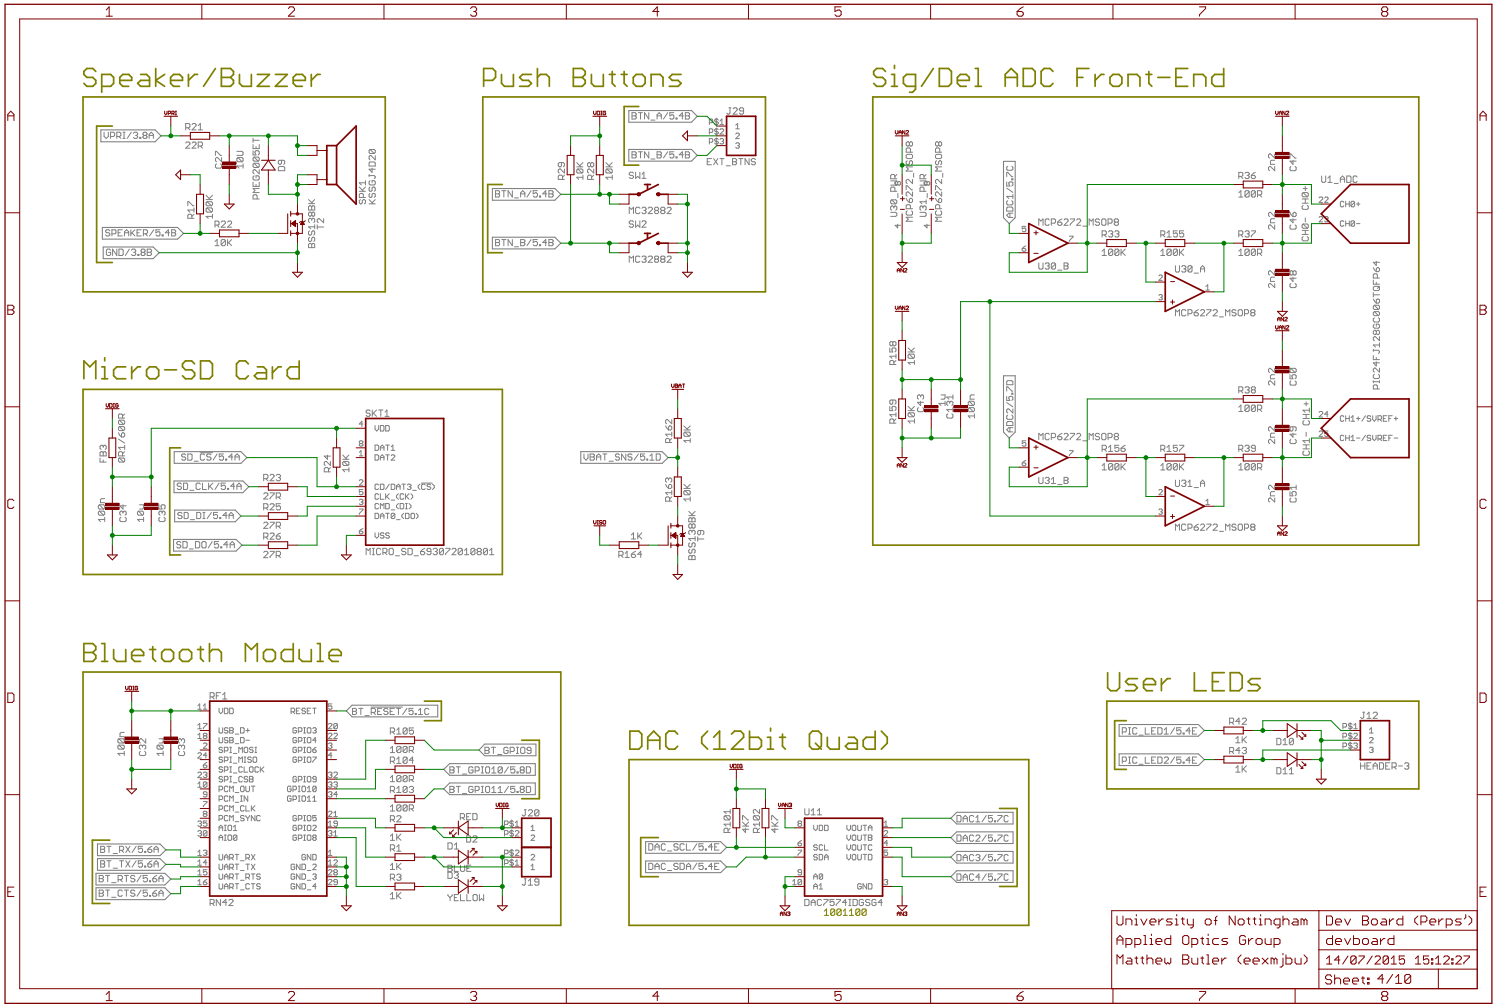

Figure 6: *Peripherals.*

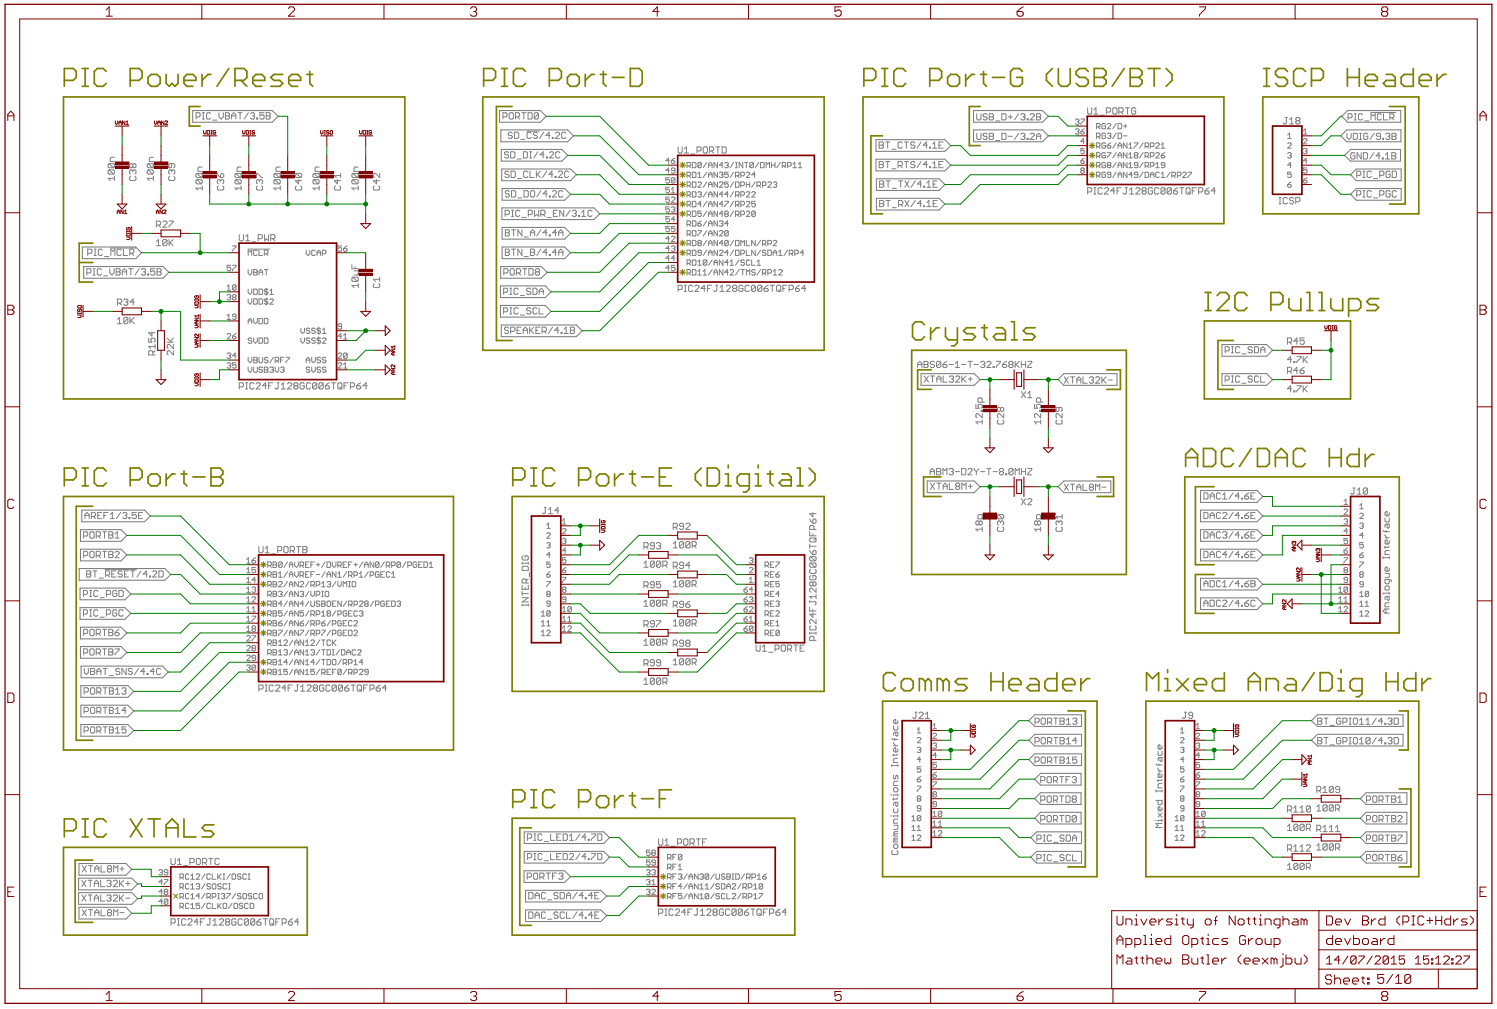

Figure 7: *Processor and headers.*

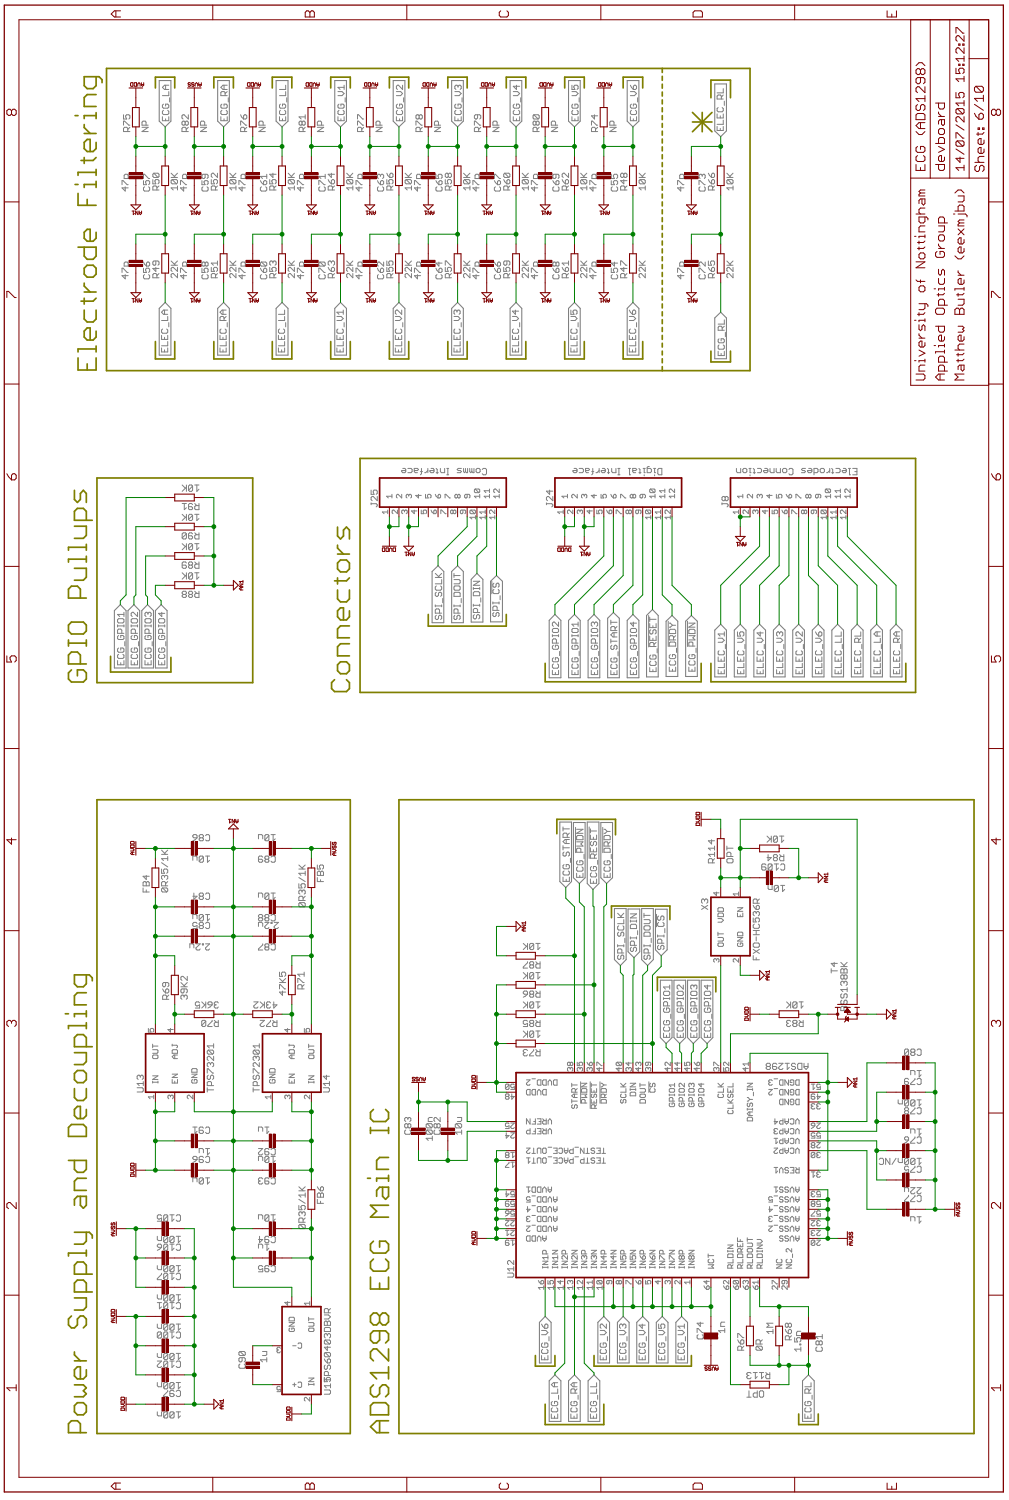

Figure 8: *ECG front-end.*

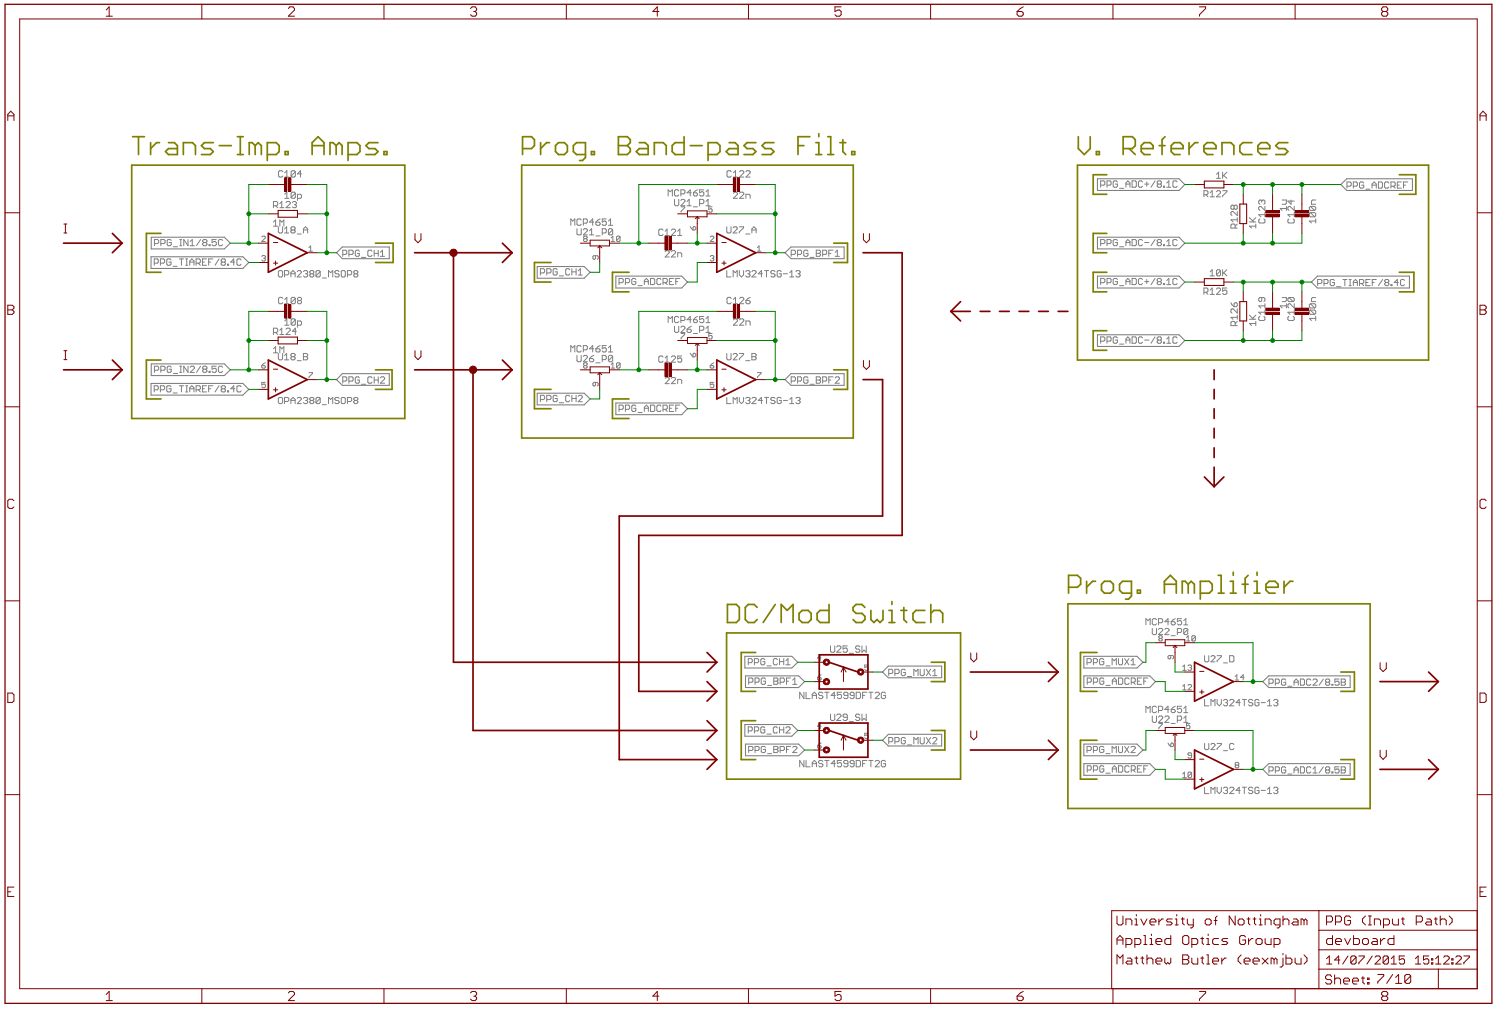

Figure 9: *PPG input paths (filter and gain stages).*

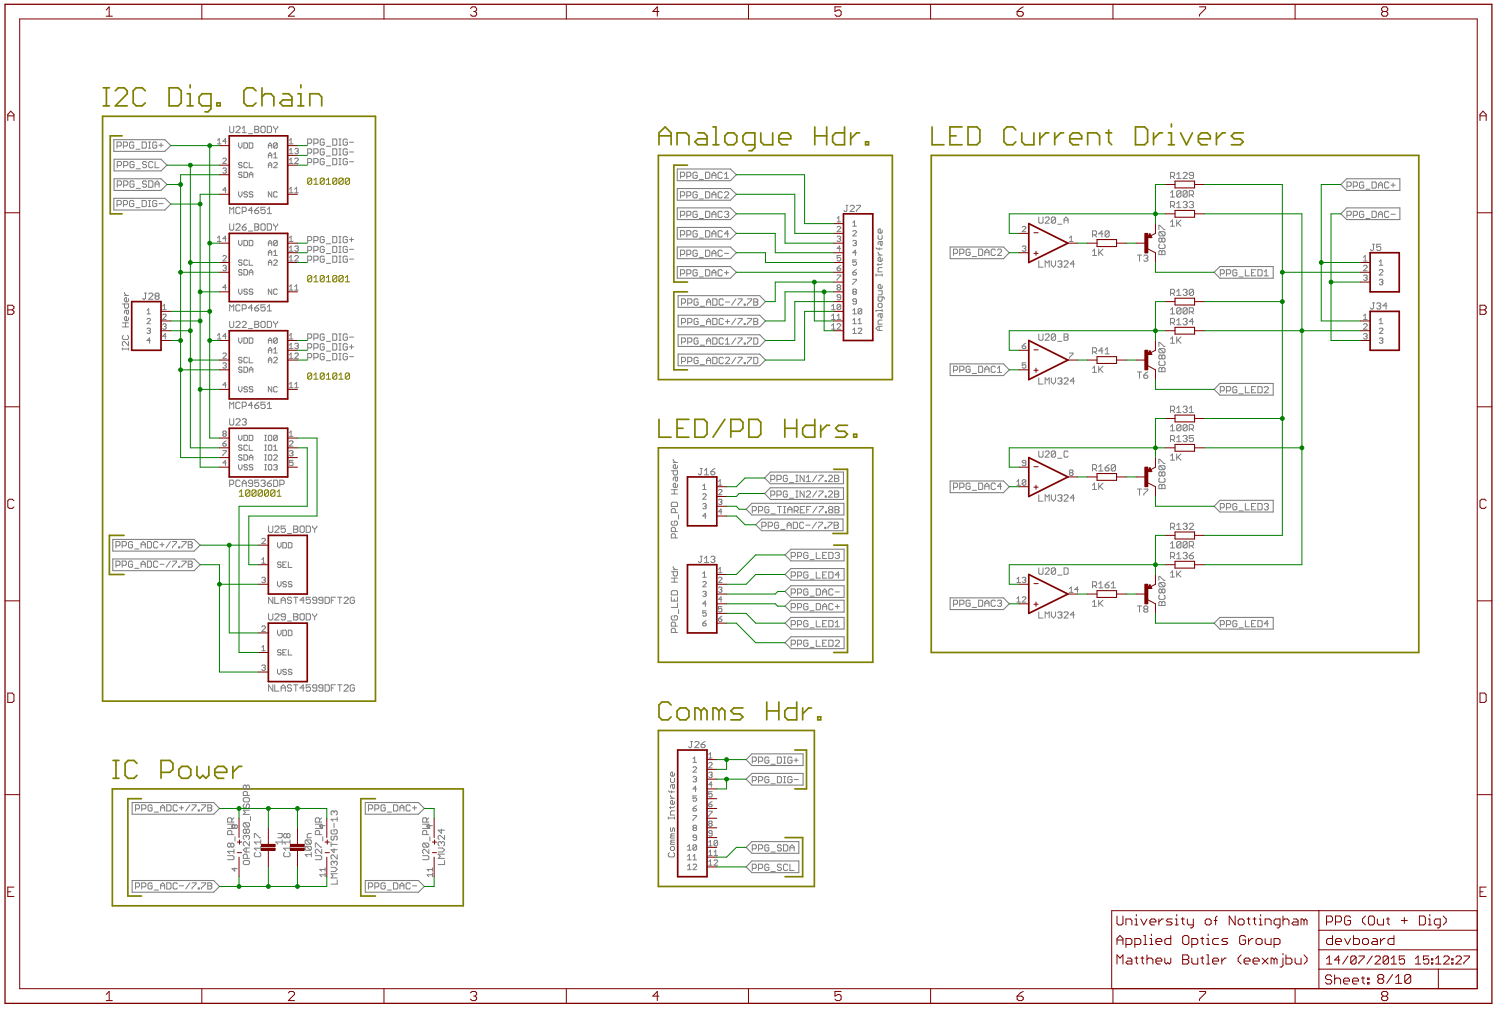

Figure 10: *PPG output paths.*

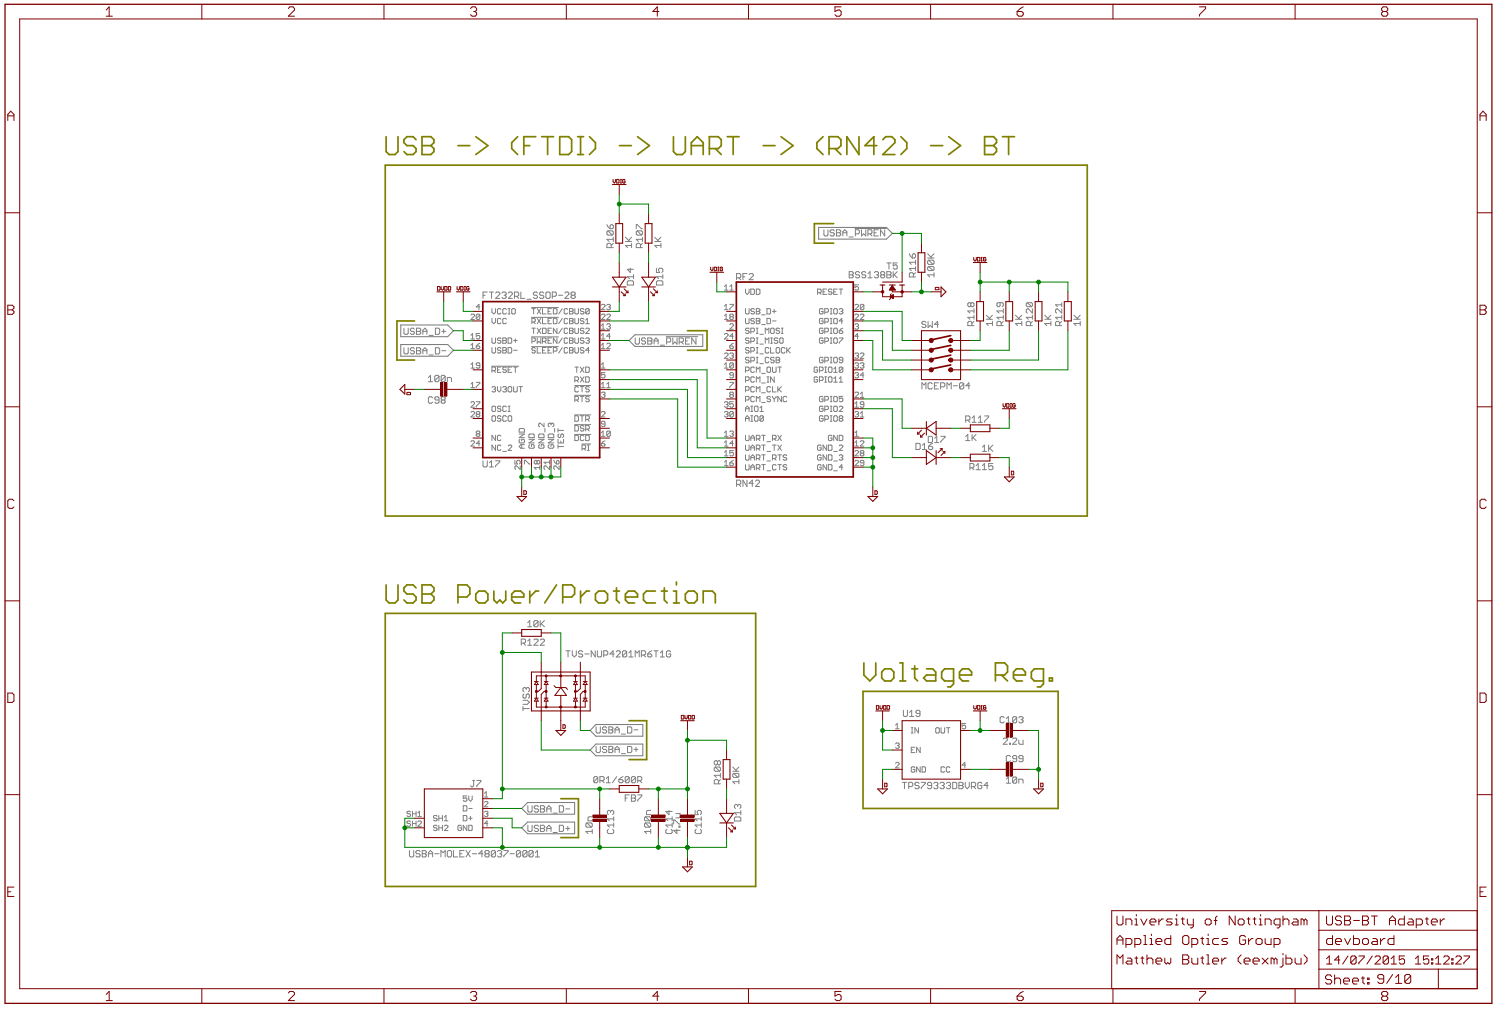

Figure 11: *USB-Bluetooth adapter.*

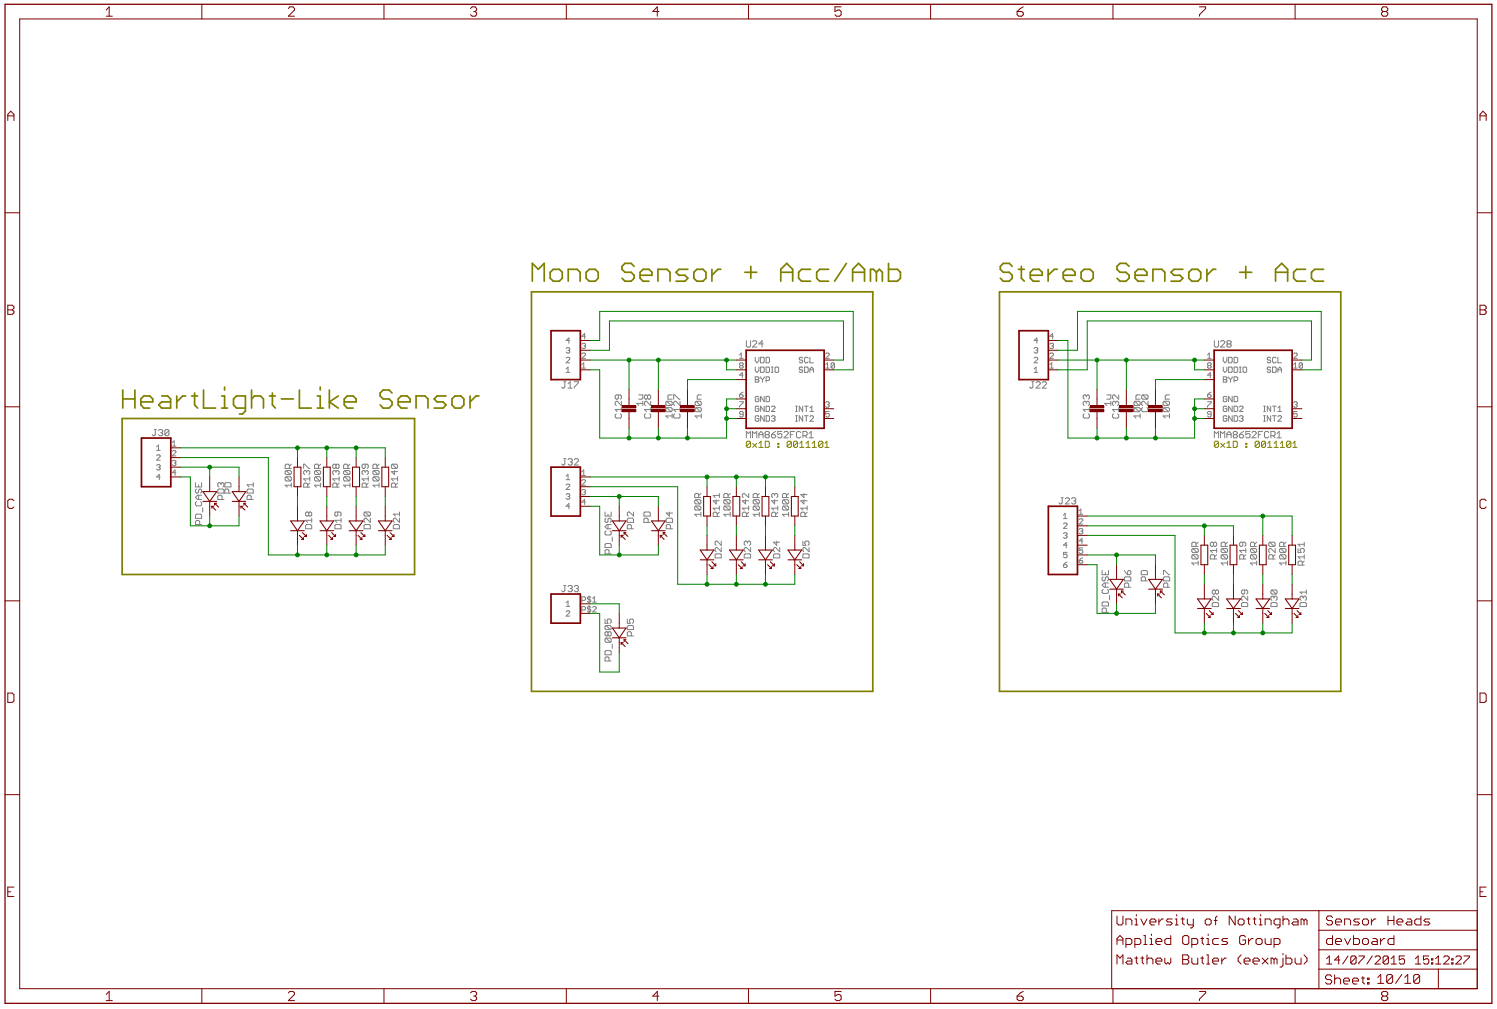

Figure 12: *Sensor heads.*

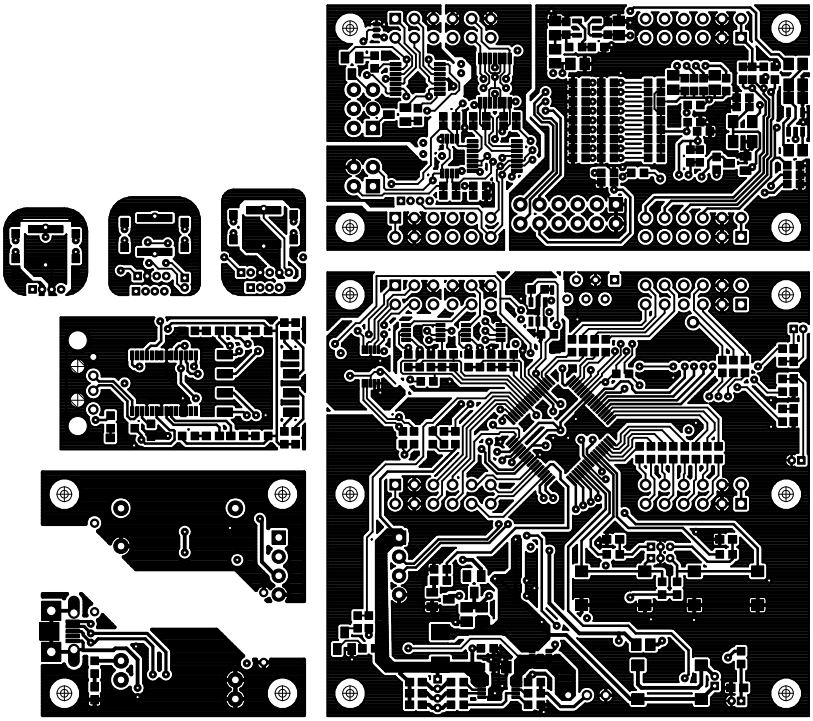

Figure 13: *Top Copper*

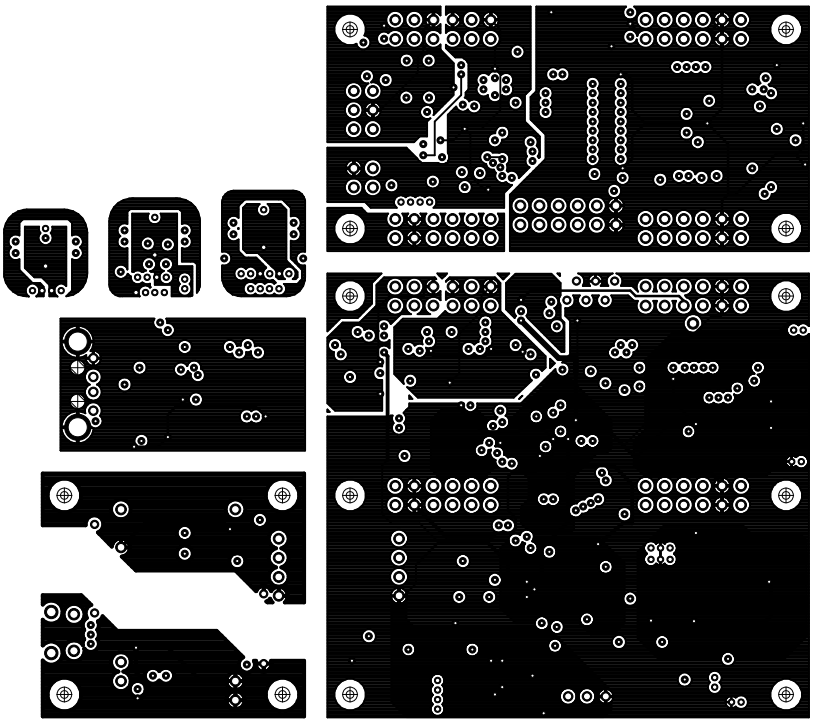

Figure 14: *Inner Layer 1*

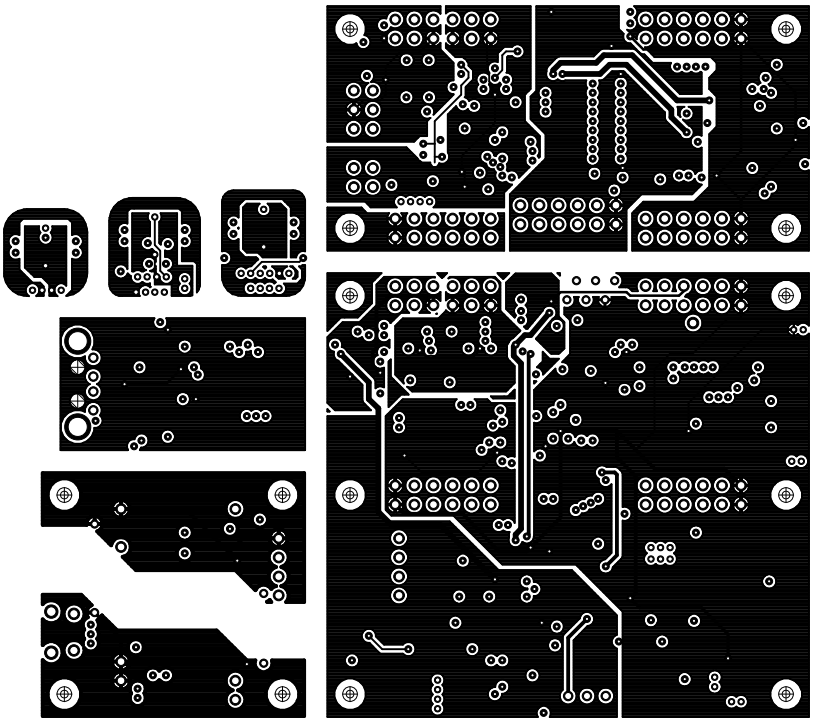

Figure 15: *Inner Layer 2*

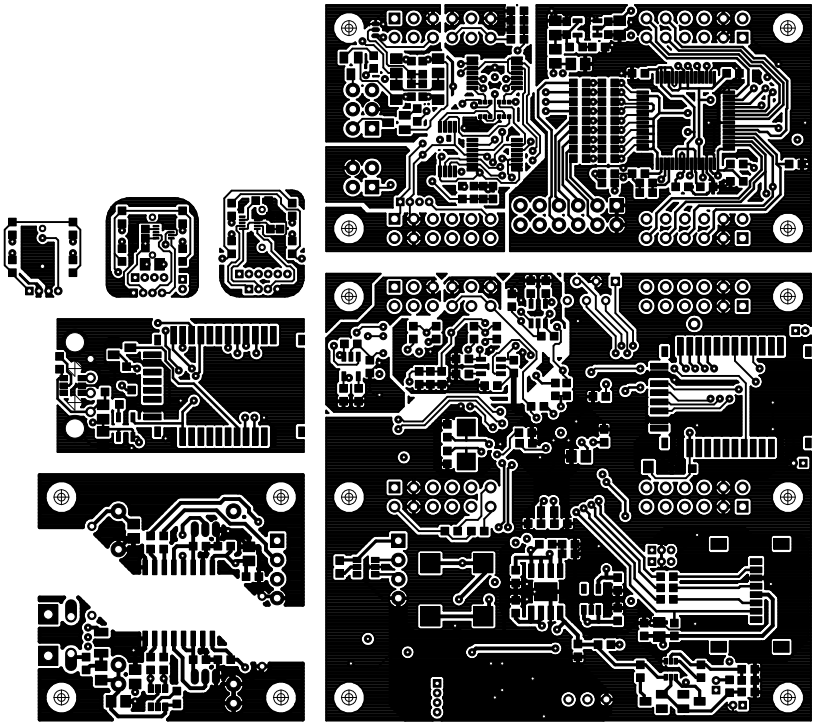

Figure 16: *Bottom Copper*

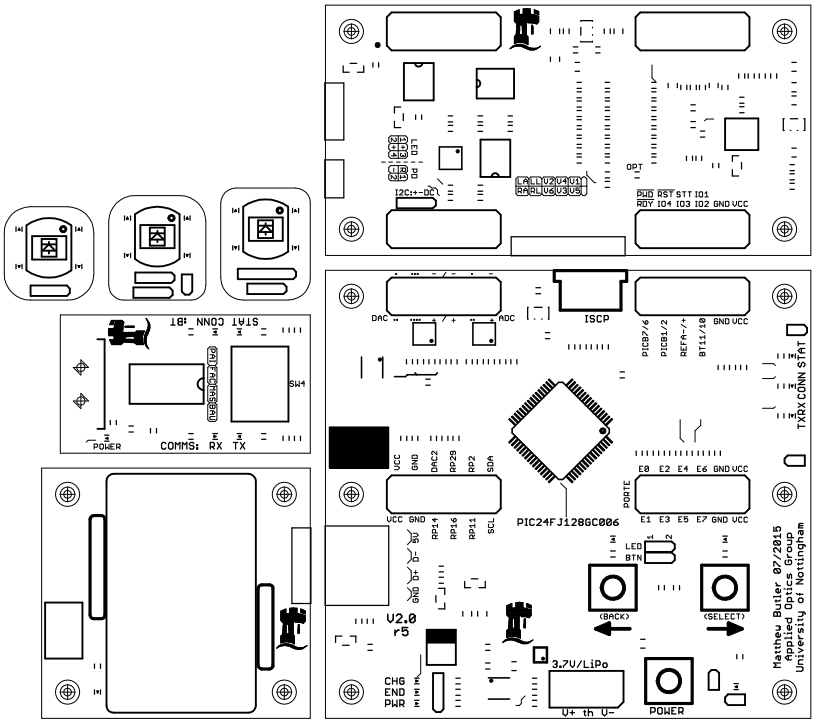

Figure 17: *Top Silk*

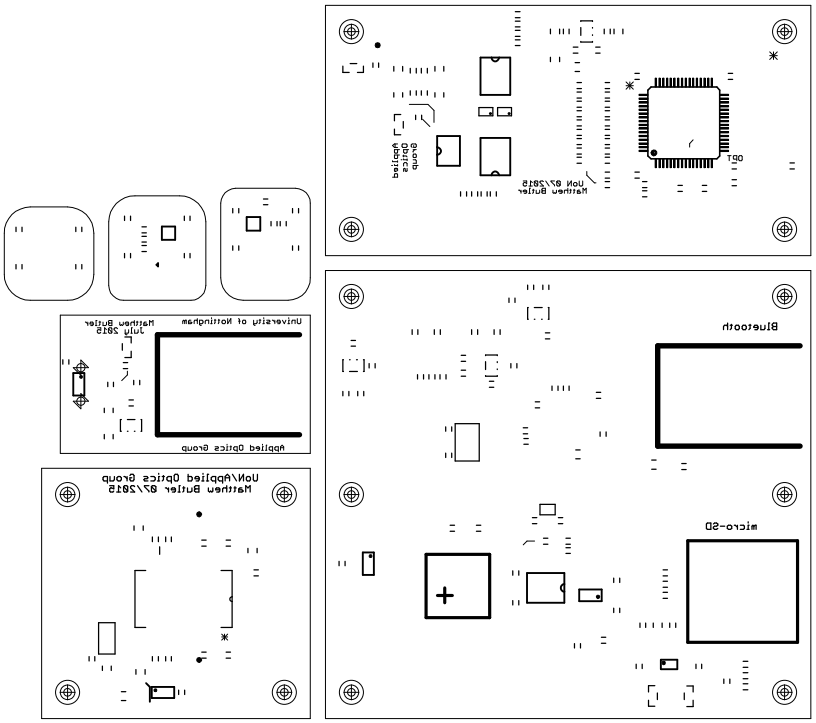

Figure 18: *Bottom Silk*

III Optical Matrix Schematic & PCB

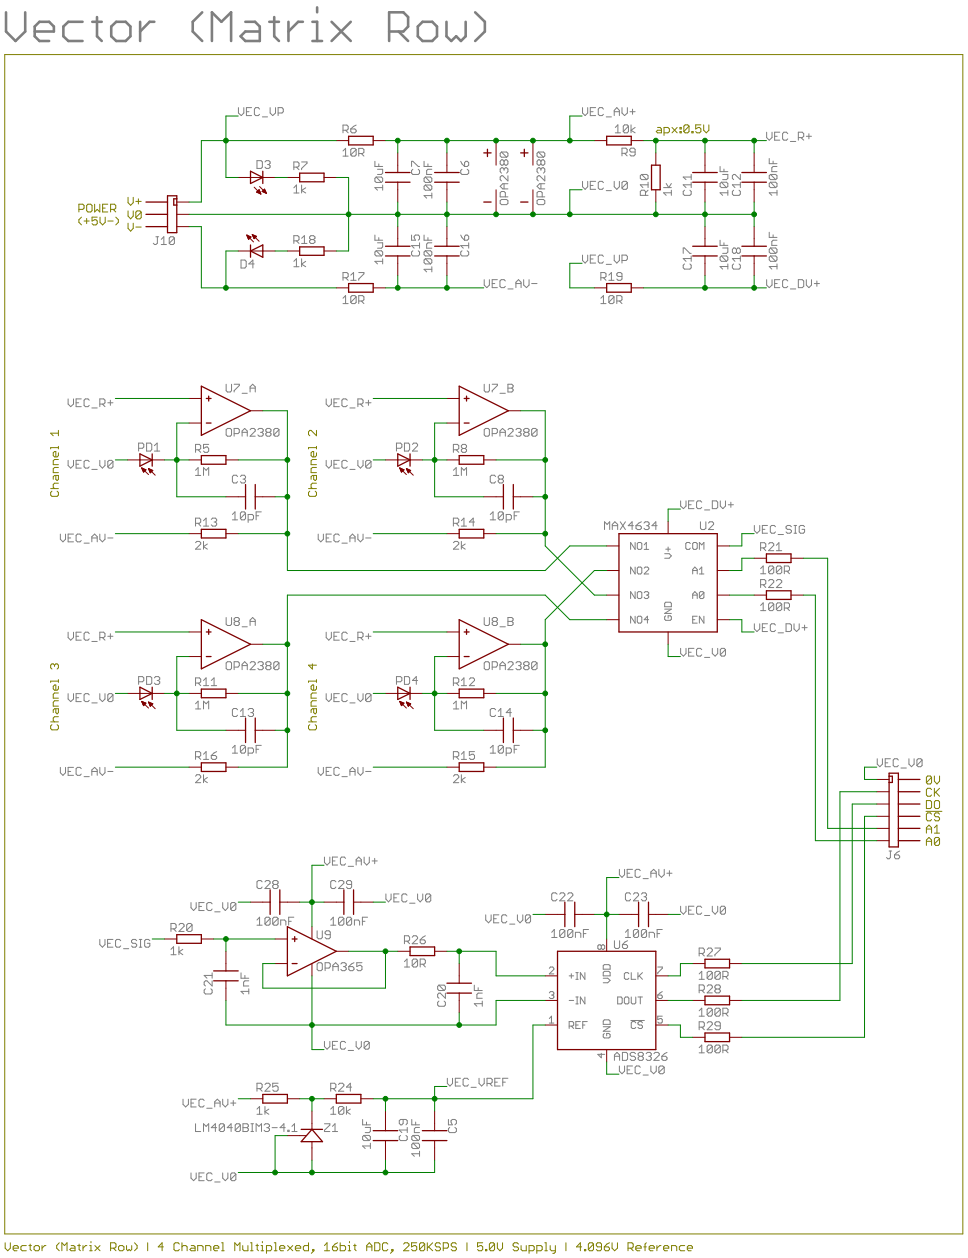

Figure 19: *Sensor inputs (amplifiers).*

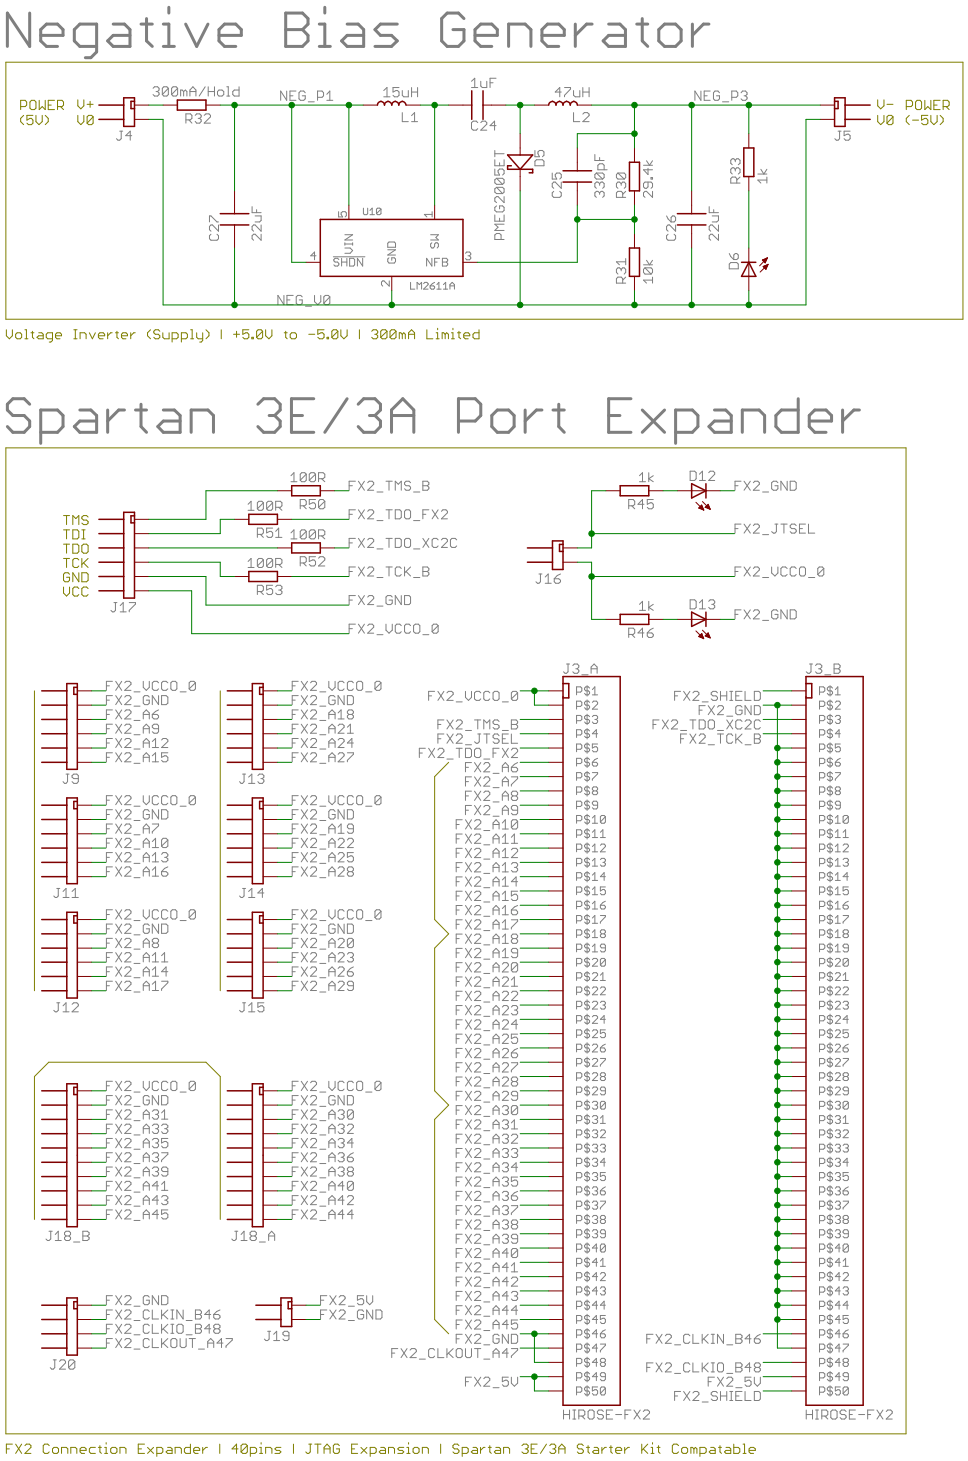

Figure 20: *FPGA connection (and negative bias generator).*

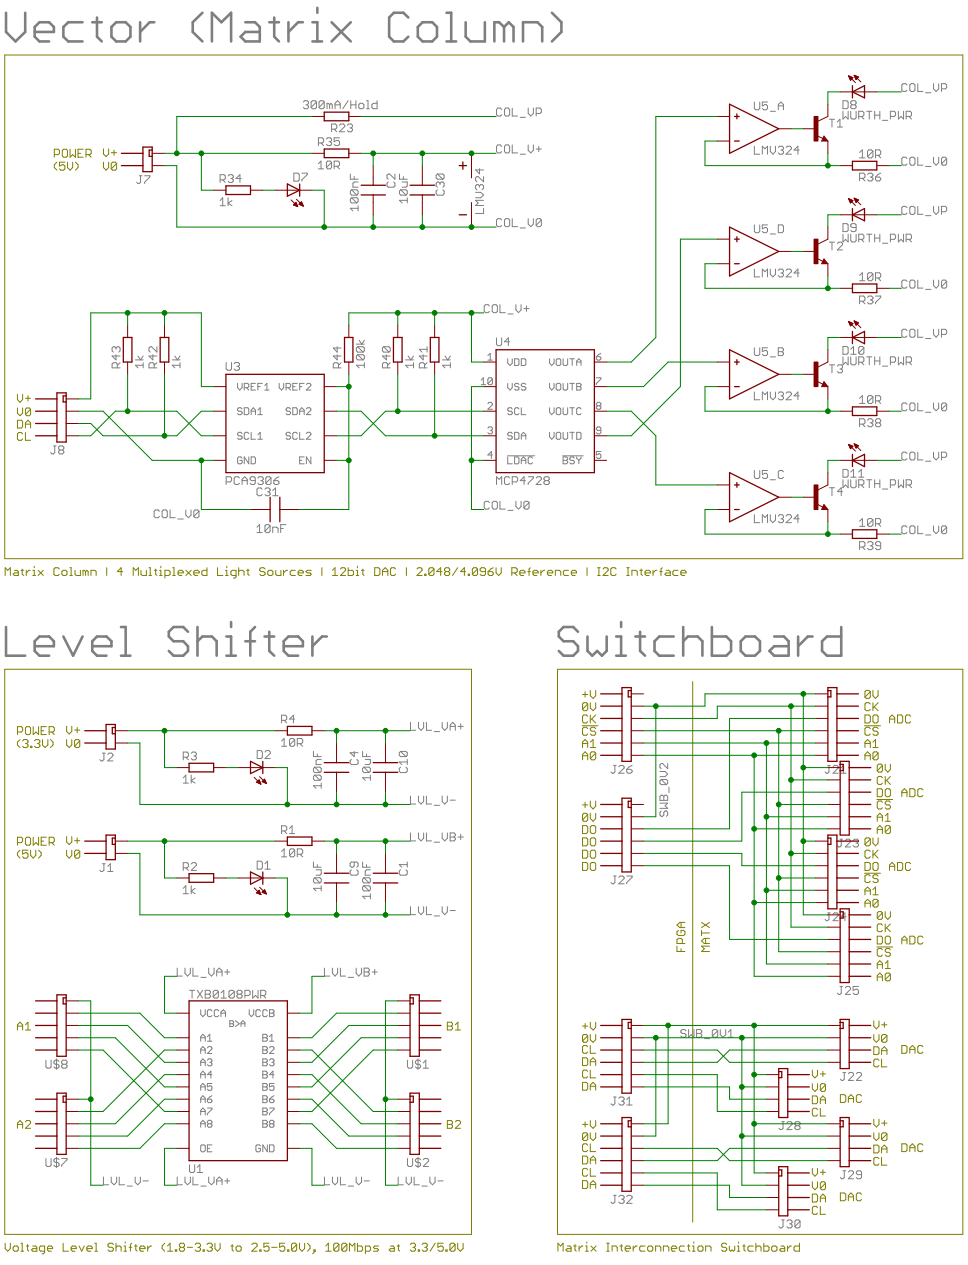

Figure 21: *LED outputs (drivers).*

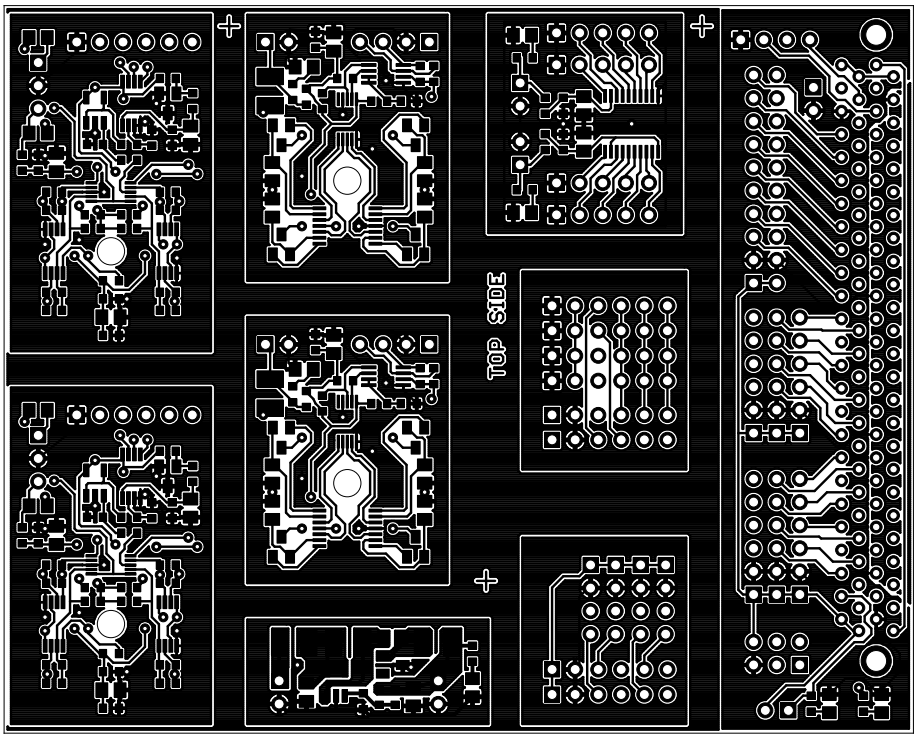

Figure 22: *Top Copper*

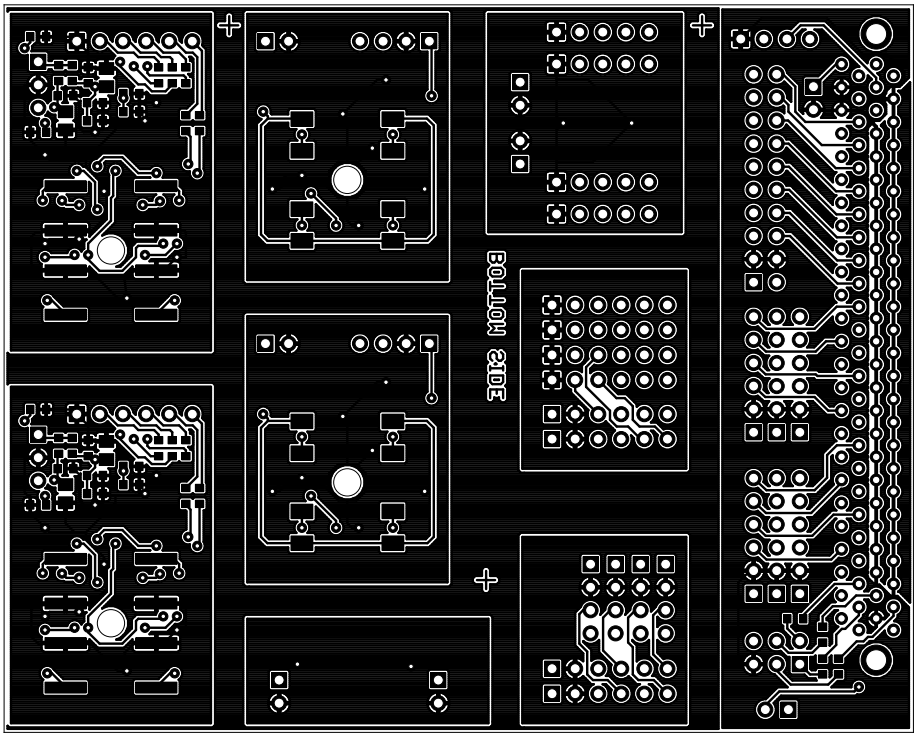

Figure 23: *Bottom Copper*

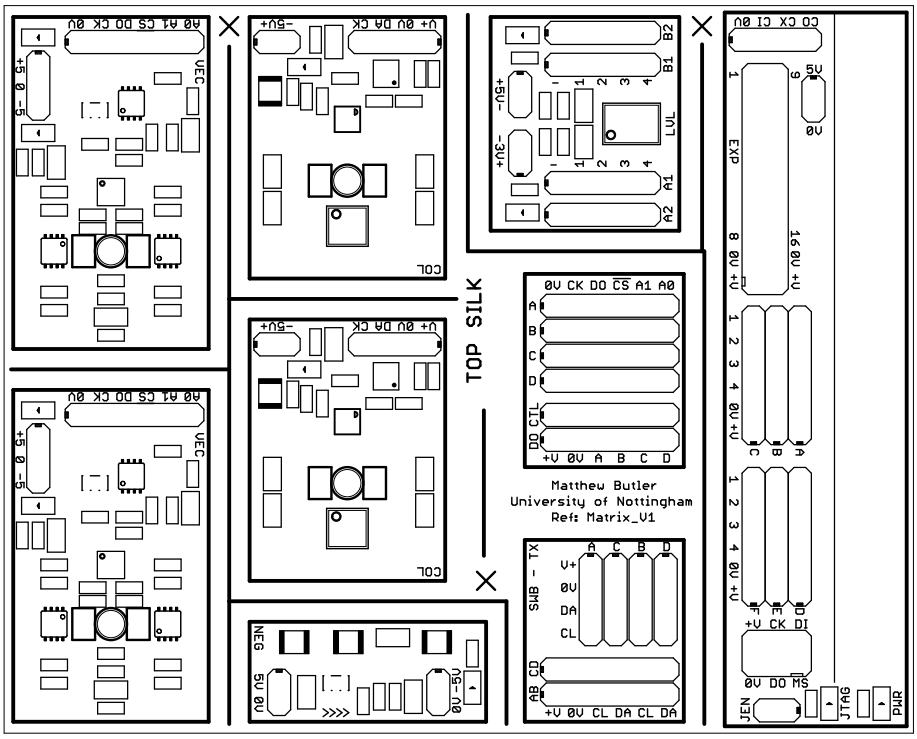

Figure 24: *Top Silk*

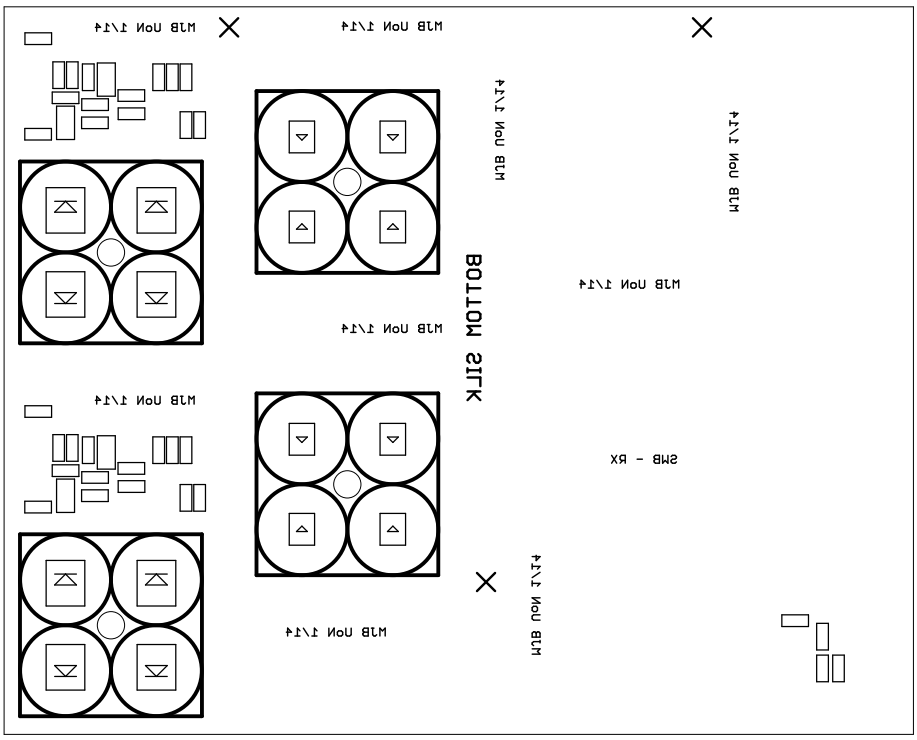

Figure 25: *Bottom Silk*

References

[1] John Allen. Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3):R1, 2007.

[2] B. C. Lacey and J. I. Lacey. Two-way communication between the heart and the brain: Significance of time within the cardiac cycle. *American Psychologist*, 33(2):99–113, 1978. doi: http://dx.doi.org/10.1037/0003-066X.33.2.99.

[3] Richard E. Klabunde. Cardiac cycle. `http://www.cvphysiology.com/Heart%20Disease/HD002.htm`, 2015. [Image; accessed 15-September-2015].

[4] Catherine Jane Owen and Jonathan Peter Wyllie. Determination of heart rate in the baby at birth. *Resuscitation*, 60(2): 213 – 217, 2004. ISSN 0300-9572. doi: https://doi.org/10.1016/j.resuscitation.2003.10.002.

[5] Hutan Ashrafian. Anatomically specific clinical examination of the carotid arterial tree. *Anatomical Science International*, 82(1):16, Mar 2007. ISSN 1447-073X. doi: 10.1111/j.1447-073X.2006.00152.x.

[6] Lokeswara Rao et al. Sajja. Role of radial artery graft in coronary artery bypass grafting. *The Annals of Thoracic Surgery*, 79(6):2180–2188, 2005.

[7] Nanfei Sun, M. Garbey, A. Merla, and I. Pavlidis. Imaging the cardiovascular pulse. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 416–421 vol. 2, June 2005. doi: 10.1109/CVPR.2005.184.

[8] American College of Cardiology. *Your Health: How to Measure Your Pulse*. Healthwise Incorporated, 2007-2015. Training Leaflet.

[9] Alexei A. Kamshilin, Ervin Nippolainen, Igor S. Sidorov, Petr V. Vasilev, Nikolai P. Erofeev, Natalia P. Podolian, and Roman V. Romashko. A new look at the essence of the imaging photoplethysmography. *Scientific Reports*, 5(10494):1–9, May 2015.

[10] William J. Federspiel and Aleksander S. Popel. A theoretical analysis of the effect of the particulate nature of blood on oxygen release in capillaries. *Microvascular Research*, 32(2):164 – 189, 1986. ISSN 0026-2862.

[11] Gordon A. Ewy, Jorge C. Rios, and Frank I. Marcus. The dicrotic arterial pulse. *Circulation*, 39:655–661, May 1969.

[12] J. Weinman, A. Hayat, and G. Raviv. Reflection photoplethysmography of arterial-blood-volume pulses. *Medical & Biological Engineering & Computing*, 15:22–31, 1977.

[13] Vincent P Crabtree and Peter R Smith. Physiological models of the human vasculature and photoplethysmography. *Electronic Systems and Control Division Research*, pages 60–63, 2003.

[14] Matthew James Hayes and Peter R. Smith. A new method for pulse oximetry possessing inherent insensitivity to artifact. *IEEE Transactions on Biomedical Engineering*, 48(4):452–461, 2001.

[15] W.G. Zijlstra, A. Buursma, and O.W. van Assendelft. *Visible and Near Infrared Absorption Spectra of Human and Animal Haemoglobin: Determination and Application*. CRC Press, 2000. ISBN 978-9067643177.

[16] P.A. Riley. Melanin. *The International Journal of Biochemistry & Cell Biology*, 29(11):1235 – 1239, 1997. ISSN 1357-2725.

[17] M. Linh Tran, Ben J. Powell, and Paul Meredith. Chemical and structural disorder in eumelanins: A possible explanation for broadband absorbance. *Biophysical Journal*, 90(3):743–752, 2006.

[18] Masimo. Masimo – single patient use – sensor range. `http://www.masimo.com/en-GB/home/signal-extraction-pulse-oximetry/masimo-set-sensors/lncs-adhesive-sensors/`. [Online; accessed 4-August-2017].

[19] Generic nellcor finger clip sensor. From `http://cardiologysupplies.com/Pulse-Oximeters/Oximeter-Sensors/Generic-Nellcor-Finger-Clip-Sensor.html`, . [Image, Accessed February 2016].

[20] Nonin earlobe clip sensor. From `http://www.farlamedical.co.uk/category_Accessories-2232/Nonin-Reusable-Ear-Clip-Sensor-For-Nonin-8500-Pulse-Oximeter/`, . [Image, Accessed March 2016].

[21] Nonin 8000r reflectance forhead sensor. From `https://www.youtube.com/watch?v=CLFguKiUp0A` thumbnail, . [Image, Accessed March 2016].

[22] D. Phan, L. Y. Siong, P. N. Pathirana, and A. Seneviratne. Smartwatch: Performance evaluation for long-term heart rate monitoring. In *2015 International Symposium on Bioelectronics and Bioinformatics (ISBB)*, pages 144–147, Oct 2015. doi: 10.1109/ISBB.2015.7344944.

[23] M. A. Hassan, A. S. Malik, N. Saad, B. Karasfi, D. Fofi, and W. Sohail. Towards health monitoring in visual surveillance. In *2016 6th International Conference on Intelligent and Advanced Systems (ICIAS)*, pages 1–6, Aug 2016. doi: 10.1109/ICIAS.2016.7824046.

[24] Logitech hd webcam c615. From `https://www.logitech.com/en-us/webcam-communications/webcams`, . [Image, Accessed March 2016].

[25] Pixelfly-vga scientific camera. From `http://www.pco.de/sensitive-cameras/pixelfly-vga/`, . [Image, Accessed March 2016].

[26] G Cennini, J Arguel, K Aksit, and A van Leest. Heart rate monitoring via remote photoplethysmography with motion artifacts reduction. *Optics Express*, 10:4867–4875, 2010.

[27] L. R. Scheker, P. G. Slattery, J. C. Firrell, and G. D. Lister. The value of the photoplethysmograph in monitoring skin closure in microsurgery. *J reconstr Microsurg*, 2(1):1–4, 1985. doi: 10.1055/s-2007-1007039.

[28] S.K. Tracy, C. Thorogood, J. Pincombe, and S. Pairman. *Midwifery: Preparation for Practice*. Elsevier Health Sciences, 2011. ISBN 9780729579285.

[29] J A Dawson, C O F Kamlin, C Wong, A B te Pas, C P F O'Donnell, S M Donath, P G Davis, and C J Morley. Oxygen saturation and heart rate during delivery room resuscitation of infants. *Archives of Disease in Childhood - Fetal and Neonatal Edition*, 94(2):F87–F91, 2009. ISSN 1359-2998. doi: 10.1136/adc.2008.141341.

[30] Nounou MI El-Amrawy F. Are currently available wearable devices for activity tracking and heart rate monitoring accurate, precise, and medically beneficial? 21(4):315–320, 2015. doi: https://doi.org/10.4258/hir.2015.21.4.315.

[31] Kelly R. Evenson, Michelle M. Goto, and Robert D. Furberg. Systematic review of the validity and reliability of consumer-wearable activity trackers. *International Journal of Behavioral Nutrition and Physical Activity*, 12(1):159, Dec 2015. ISSN 1479-5868. doi: 10.1186/s12966-015-0314-1.

[32] Alrick B. Hertzman. Photoelectric plethysmography of the fingers and toes in man. *Experimental Biology and Medicine*, 37(3):529–534, 1937.

[33] P. J. Hanzlik, Floyd DeEds, and B. Terada. A simple method of demonstrating changes in blood supply of the ear and effects of some measures. *Journal of Pharmacology and Experimental Therapeutics*, 56(2):194–204, 1936.

[34] Hans Molitor and Michael Kniazuk. A new bloodless method for continuous recording of peripheral circulatory changes. *Journal of Pharmacology and Experimental Therapeutics*, 57(1):6–18, 1936.

[35] Tor O Klemsdal, Tomas L G Andersson, Jørgen Matz, Gordon A A Ferns, Knut Gjesdal, and Erik E Änggård. Vitamin e restores endothelium dependent vasodilatation in cholesterol fed rabbits: in vivo measurements by photoplethysmography. *Cardiovascular Research*, 28:1397–1402, 1994.

[36] Claude Julien. The enigma of mayer waves: Facts and models. *Cardiovascular Research*, 70:12–21, 2006.

[37] Kenneth E Nelson, Nicette Sergueef, Celia M Lipinski, Arina R Chapman, and Thomas Glonek. Cranial rhythmic impulse related to the traube-hering-mayer oscillation: comparing laser-doppler flowmetry and palpation. *The Journal of the American Osteopathic Association*, 101(3):163–173, March 2001.

[38] Ben J A Janssen, Simon C Malpas, Sandra L Burke, and Geoffrey A Head. Frequency-dependent modulation of renal blood flow by renal nerve activity in conscious rabbits. *American Journal of Physiology*, 273(2:2):597–608, 1997.

[39] Simon C Malpas and Don E Burgess. Renal sna as the primary mediator of slow oscillations in blood pressure during hemorrhage. *American Journal of Physiology - Heart and Circulatory Physiology*, 279(3): 1299–1306, 2000.

[40] Wearable technologies. `http://www.wearable-technologies.com`. Accessed: 27/11/2015.

[41] Brian Ballard. Tracing the roots of wearable technology. `http://techcrunch.com/2015/11/24/tracing-the-roots-of-wearable-technology/`, November 2015. Crunch Network, Accessed: 27/11/2015.

[42] Sharon Profis. Do wristband heart trackers actually work? a checkup. `http://www.cnet.com/uk/news/how-accurate-are-wristband-heart-rate-monitors/`, May 2014. CNET, Accessed: 27/11/2015.

[43] Shane Richmond. The real world wrist-based heart rate monitor test: Are they accurate enough? `http://www.wareable.com/fitness-trackers/heart-rate-monitor-accurate-comparison-wrist`, July 2015. Accessed: 27/11/2015.

[44] KG. Voogdt, AC. Morrison, FE. Wood, RM. van Elburg, and JP. Wyllie. A randomised, simulated study assessing auscultation of heart rate at birth. *Resuscitation*, 81(8):1000–1003, 2010.

[45] Perlman JM, Wyllie J, Kattwinkel J, Atkins DL, Chameides L, Goldsmith JP, Guinsburg R, Hazinski MF, Morley C, Richmond S, Simon WM, Singhal N, Szyld E, Tamura M, and Velaphi S. Part 11: neonatal resuscitation: 2010 international consensus on cardiopulmonary resuscitation and emergency cardiovascular care science with treatment recommendations. *Circulation*, 122:516–538, 2010. supplement 2.

[46] Nicola J Robertson. Air or 100% oxygen for asphyxiated babies? Time to decide. *Critical Care*, 9(2):128–130, 2005.

[47] Bethanie Christopher. *NRP Study Guide*. Lifesaver Educational, January 2012. http://www.lifesavered.com/cmsfiles/nrp-guide-re.pdf Accessed: 30/11/2015.

[48] Virginia Apgar. A proposal for a new method of evaluation of the newborn infant. *Current Researches in Anesthesia and Analgesia*, 32 (4):260–267, 1953.

[49] Charlotte Howell, Kate Grady, and Charles Cox. *Managing Obstetric Emergencies and Trauma: The MOET Course Manual*, chapter 8. Cambridge University Press, 3rd edition, May 2014.

[50] New heart rate monitor for babies who need resuscitation at birth. https://www.action.org.uk/our-research/new-heart-rate-monitor-babies-who-need-resuscitation-birth, August 2007. Action Medical Research (for children) Accessed: 30/11/2015.

[51] A green light for improved resuscitation. Midwives Magazine (also at https://www.rcm.org.uk/news-views-and-analysis/analysis/a-green-light-for-improved-resuscitation), August 2010. Accessed: 30/11/2015.

[52] Kevin K. Tremper. Pulse oximetry. *Chest*, 95(4):713–715, 1989.

[53] T.L. Rusch, R. Sankar, and J.E. Scharf. Signal processing methods for pulse oximetry. *Computers in Biology and Medicine*, 26(2):143 – 159, 1996. ISSN 0010-4825. doi: https://doi.org/10.1016/0010-4825(95)00049-6.

[54] John A. Crowe and Damianos Damianou. The wavelength dependence of the photoplethysmogram and its implication to pulse oximetry. pages 2423–2424, December 1992. doi: 10.1109/IEMBS.1992.592733.

[55] Diwei He, Stephen Morgan, Dimitrios Trachanis, Jan Van Hese, Dimitrios Drogoudis, Franco Fummi, Francesco Stefanni, Valerio Guarnieri, and Barrie Hayes-Gill. A single-chip cmos pulse oximeter with on-chip lock-in detection. *Sensors*, 15(7):17076–17088, June 2015.

[56] Lepeschkin EE. Practical evaluation of the electrocardiogram. *Archives of Internal Medicine*, 126(6):1074–1075, 1970. doi: 10.1001/archinte.1970.00310120136020.

[57] Domenico Corrado, Antonio Pelliccia, Hein Heidbuchel, Sanjay Sharma, Mark Link, Cristina Basso, Alessandro Biffi, Gianfranco Buja, Pietro Delise, Ihor Gussac, Aris Anastasakis, Mats Borjesson, Hans Halvor Bjrnstad, Franois Carr, Asterios Deligiannis, Dorian Dugmore, Robert Fagard, Jan Hoogsteen, Klaus P. Mellwig, Nicole Panhuyzen-Goedkoop, Erik Solberg, Luc Vanhees, Jonathan Drezner, N.A. Mark Estes, III, Sabino Iliceto, Barry J. Maron, Roberto Peidro, Peter J. Schwartz, Ricardo Stein, Gaetano Thiene, Paolo Zeppilli, and

William J. McKenna. Recommendations for interpretation of 12-lead electrocardiogram in the athlete. *European Heart Journal*, 31(2):243–259, 2010. doi: 10.1093/eurheartj/ehp473.

[58] Toshiyo Tamura, Yuka Maeda, Masaki Sekine, and Masaki Yoshida. Wearable photoplethysmographic sensorspast and present. *Electronics*, 3(2):282–302, 2014. ISSN 2079-9292. doi: 10.3390/electronics3020282.

[59] *Application Note 291565 Comparison of EPIC electrodes with Ag/AgCl for ECG measurement*. Plessey Semiconductors Ltd., Design & Technology Centre, Delta 500, Delta Business Park, Great Western Way, Swindon, UK SN5 7XE, . Application Note, Issue 1.

[60] *Application Note 291566 Non-contact ECG measurement using EPIC*. Plessey Semiconductors Ltd., Design & Technology Centre, Delta 500, Delta Business Park, Great Western Way, Swindon, UK SN5 7XE, . Application Note, Issue 1.

[61] Thomas J. Sullivan, Stephen R. Deiss, and Gert Cauwenberghs. A low-noise, non-contact eeg/ecg sensor. In *Biomedical Circuits and Systems Conference, BIOCAS*, pages 154–157. IEEE, November 2007.

[62] Ming-Zher Poh, Daniel J. McDuff, and Rosalind W. Picard. Non-contact, automated cardiac pulse measurements using video imaging and blind source separation. *Optics Express*, 18(10):10762–10774, 2010.

[63] Alexei A. Kamshilin, Serguei Miridonov, Victor Teplov, Riku Saarenheimo, and Ervin Nippolainen. Photoplethysmographic imaging of high spatial resolution. *Biomedical Optics Express*, 2(4):996–1006, 2011.

[64] Wim Verkruysse, Lars O Svaasand, and J Stuart Nelson. Remote plethysmographic imaging using ambient light. *Optics Express*, 16 (26):21434–21445, 2008.

[65] Xiaoqin Luo, Fen Miao, and Ye Li. Ppg and ecg feature comparison for healthy people and hypertensive patients. pages 785–787, Hong Kong and Shenzhen, China, January 2012. IEEE-EMBS International Conference on Biomedical and Health Informatics.

[66] M. Bolanos, H. Nazeran, and E. Haltiwanger. Comparison of heart rate variability signal features derived from electrocardiography and photoplethysmography in healthy individuals. pages 4289–4294, New York City, USA, August & September 2006. 28th IEEE EMBS Annual International Conference.

[67] G. Lu, F. Yang, J. A. Taylor, and J. F. Stein. A comparison of photoplethysmography and ecg recording to analyse heart rate variability in healthy subjects. *Journal of Medical Engineering & Technology*, 33 (8):634–641, 2009.

[68] Victor H Frankel and Kenneth J McLeod. Calf muscle pump stimulation as an adjunct to orthopaedic surgery. *Surgical Technology International*, 14:297–304, 2005.

[69] M J Butler, J A Crowe, B R Hayes-Gill, and P I Rodmell. Motion limitations of non-contact photoplethysmography due to the optical and topological properties of skin. *Physiological Measurement*, 37(5): N27–37, April 2016.

[70] Levi Benjamin Wood. *Motion Artifact Reduction for Wearable Photoplethysmogram Sensors Using Micro Accelerometers and Laguerre Series Adaptive Filters*. Phd thesis, Massachusetts Institute of Technology, June 2008.

[71] (Unknown). *Motion artifact reduction with active noise cancellation and accelerometer for wearable Photoplethysmography*. Title page not fully translated, (Unknown location), December 2006.

[72] Boreom Lee, Jonghee Han, Hyun Jae Baek, Jae Hyuk Shin, Kwang Suk Park, and Won Jin Yi. Improved elimination of motion artifacts from a photoplethysmographic signal using a kalman smoother with simultaneous accelerometry. *Physiological Measurement*, 31:1585–1603, 2010.

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

[74] Bennett Ames Fallow. *Influence of Skin Type and wavelength on Light Wave Reflectance*. Msc. thesis, The University of Texas at Austin, May 2012.

[75] J. Lee, K. Matsumura, K. i. Yamakoshi, P. Rolfe, S. Tanaka, and T. Yamakoshi. Comparison between red, green and blue light reflection photoplethysmography for heart rate monitoring during motion. In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 1724–1727, July 2013. doi: 10.1109/EMBC.2013.6609852.

[76] J. A. C. Patterson, D. C. McIlwraith, and G. Z. Yang. A flexible, low noise reflective ppg sensor platform for ear-worn heart rate monitoring. In *2009 Sixth International Workshop on Wearable and Implantable Body Sensor Networks*, pages 286–291, June 2009. doi: 10.1109/BSN.2009.16.

[77] Intertek. Iec/en 62471 for led lighting products. Standards information leaflet, Intertek Group plc.

[78] CELMA. Optical safety of led lighting. Safety Briefing 1, European Lamp Companies Federation, July 2011.

[79] Sangeeta Bagha and Laxmi Shaw. A real time analysis of ppg signal for measurement of spo2 and pulse rate. *International Journal of Computer Applications*, 36(11):45–50, 2010.

[80] Kent H. Lundberg. Noise sources in bulk cmos. MIT Online Library, 2002.

[81] Susannah Fleming, Matthew Thompson, Richard Stevens, Carl Heneghan, Annette Plüddemann, Ian Maconochie, Lionel Tarassenko, and David Mant. Normal ranges of heart rate and respiratory rate in children from birth to 18 years: a systematic review of observational studies. *Lancet*, 377(9770):1011–1018, March 2011.

[82] Kenneth Gerard Humphreys. *An Investigation of Remote Non-Contact Photoplethysmography and Pulse Oximetry*. PhD thesis, Department of Electronic Engineering Faculty of Science and Engineering, NUI Maynooth, August 2007.

[83] J M Bland and D G Altman. Statistical methods for assessing agreement between two methods of clinical measurement. *Lancet*, 327 (8476):307–310, 1986.

[84] C Orphanidou, T Bonnici, P Charlton, D Clifton, D Vallance, and L Tarassenko. Signal-quality indices for the electrocardiogram and photoplethysmogram: derivation and applications to wireless monitoring. *IEEE Journal of Biomedical and Health Informatics*, 19(3): 832–838, May 2015.

[85] Zhilin Zhang. Photoplethysmography-based heart rate monitoring in physical activities via joint sparse spectrum reconstruction. *IEEE Transactions on Biomedical Engineering*, 62(8):1902–1910, 2015.

[86] Zhe Lin, Jin Zhang, Yanjiao Chen, and Qian Zhang. Heart rate estimation using wrist-acquired photoplethysmography under different types of daily life motion artifact. In *IEEE International Conference on Communications*, Communications for E-Health, pages 2095–2100, London, UK, June 2015.

[87] Sonnia María López-Silva, Romano Giannetti, María Luisa Dotor, Juan Pedro Silveira, Dolores Golmayo, Francisco Miguel-Tobal, Amaya Bilbao, Mercedes Galindo, and Pilar Martn-Escudero. Heuristic algorithm for photoplethysmographic heart rate tracking during maximal exercise test. *Journal of Medical and Biological Engineering*, 32(3):181–188, 2012.

[88] U. Grenander. *Probability and Statistics: The Harald Cramér Volume*. Wiley Publications in Statistics. Alqvist & Wiksell, 1959.

[89] Eduardo E. Benarroch, Tonette L. Opfer-Gehrking, and Phillip A. Low. Use of the photoplethysmographic technique to analyze the valsalva maneuver in normal man. *Muscle & Nerve*, 14:1165–1172, December 1991.

[90] E. Kviesis-Kipge, J. Zaharans, O. Rubenis, and A. Grabovskis. A photoplethysmography device for multipurpose blood circulatory system assessment. In *Novel Biophotonic Techniques and Applications*. European Conferences on Biomedical Optics, June 2011.

[91] Chakravarti, Laha, and Roy. *Handbook of Methods of Applied Statistics*, volume 1, pages 392–394. John Wiley and Sons, 1967.

[92] Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945. ISSN 00994987.

[93] Aapo Hyvärinen and Erkki Oja. Independent component analysis: Algorithms and applications. *Neural Networks*, 13(4-5):411–430, 2000.

[94] Maik Pflugradt, Marcus Rose, and Reinhold Orglmeister. A novel method for motion artifact removal in wearable ppg sensors based on blind source separation. *Biomedical Engineering (Technology)*, 58: Suppl. 1, 2013.

[95] Fulai Peng, Zhengbo Zhang, Xiaoming Gou, Hongyun Liu, and Weidong Wang. Motion artifact removal from photoplethysmographic signals by combining temporally constrained independent component analysis and adaptive filter. *BioMedical Engineering OnLine*, 13(50), 2014.

[96] Meir Nitzan, Ayal Romem, and Robert Koppel. Pulse oximetry: fundamentals and technology update. *Medical Devices: Evidence and Research*, 7:231–239, July 2014.

[97] Fabrizio Esposito, Erich Seifritz, Elia Formisano, Renato Morrone, Tommaso Scarabino, Gioacchino Tedeschi, Sossio Cirillo, Rainer Goebel, and Francesco Di Salle. Real-time independent component analysis of fmri time-series. *NeuroImage*, 20:2209–2224, August 2003.