

Low-cost Non-contact Monitoring of Small Motions in the Human Body

University of Cape Town



Kukhokuhle Tsengwa

Supervised by:

A/Prof. Amir Patel,

A/Prof. Fred Nicolls,

Dr. Stephen Paine

October 2023

Declaration

I know the meaning of plagiarism and declare that all the work in the document, save for that which is properly acknowledged, is my own. This dissertation has been submitted to the Turnitin module (or equivalent similarity and originality checking software) and I confirm that my supervisor has seen my report and any concerns revealed by such have been resolved with my supervisor.

Name: Kukhokuhle Tsengwa
Signature: K. Tsengwa
Date: 02 October 2023

Acknowledgements

This work could not have been completed, carried out or even begun without the assistance of my supervisors A/Prof. Amir Patel, A/Prof. Fred Nicolls and Dr. Stephen Paine. I appreciate too the support I received from A/Prof. Yumna Albertus. Thank you, your mentorship is highly appreciated.

In the African Robotics Unit, where I spent most of my days, it always felt like home. I would like to extend my sincere gratitude to the people who made this so. To all my colleagues who I worked and had fun with in the lab – thank you for everything.

To my mom, brother and sisters, your love and support is not lost on me. Thank you.

Abstract

In this dissertation, we discuss how Frequency Modulated Continuous Wave (FMCW) millimetre wave (mmWave) radar can be used to monitor small motions in the human body. We demonstrate our technique on three applications: heart rate monitoring, respiration rate monitoring and muscle activity monitoring. On all three applications, we achieve performance comparable to the current state of the art. For respiration rate and muscle activity monitoring, we also demonstrate that the radar can be substituted for ordinary cameras.

Our muscle activity monitoring approach relies on the fact that when muscles contract, they change in length while their volume remains constant. This leads to changes in the thickness of the muscle. Our approach measures these thickness changes without contact. Sonomyography (SMG) uses ultrasound to measure the same signal but with contact. To the best of the author's knowledge, this work is the first to demonstrate the monitoring of this signal non-contactly.

The motion of a target within the field of view of a radar can be monitored by range estimation. However, if the motion is of a sufficiently small magnitude, changes in the position of the target do not lead to changes in the estimated range. With FMCW mmWave radar, motion can be measured with micron-level accuracy by extracting the phase signal. By performing frequency analysis on this phase signal, human heart and respiration rates were estimated with 98.5% and 99.4% accuracy, respectively. The phase signal was also used to monitor the muscle activity of the vastus lateralis muscle, achieving results very similar to other approaches that measure muscle activity through muscle structural changes. It was also shown that two cameras can be used to mimic the radar phase response to monitor respiration rate and muscle activity with comparable performance.

Contents

1	Introduction	8
1.1	Background	8
1.2	Objectives and Scope	8
1.3	Outcomes	9
1.4	Overview	9
1.4.1	Literature Review	10
1.4.2	Frequency Modulated Continuous Wave (FMCW) Radar Theory	10
1.4.3	Heart Rate Estimation with Radar	10
1.4.4	Respiration Rate Estimation with Radar and Cameras	11
1.4.5	Muscle Activity Monitoring with Radar and Cameras	11
1.4.6	Limitations, Recommendations and Conclusions	11
2	Literature Review	12
2.1	Introduction	12
2.2	Heart Rate	12
2.2.1	Electrocardiography (ECG)	13
2.2.2	Non-contact monitoring of heart rate	14
2.3	Respiration Rate	14
2.3.1	Other respiration rate monitoring approaches	15
2.3.2	Non-contact monitoring of respiration rate	15
2.4	Muscle Activity	16
2.4.1	Electromyography (EMG)	16
2.4.2	Non-contact monitoring of muscle activity	17
2.4.3	Summary	18
3	Frequency Modulated Continuous Wave (FMCW) Radar Theory	19
3.1	Introduction	19
3.2	FMCW Radar Data	21
3.3	Small Motion Model	21
3.4	Received Echo Model	22
3.5	Range Estimation	24
3.6	Phase Analysis	26
4	Heart Rate Estimation with Radar	32
4.1	Introduction	32
4.2	Experimental Design	32
4.3	Results	33
5	Respiration Rate Estimation with Radar and Cameras	38
5.1	Introduction	38
5.2	Camera Based Phase Estimation	38
5.3	Experimental Design	42
5.4	Results	42

6	Muscle Activity Monitoring with Radar and Cameras	50
6.1	Introduction	50
6.2	Experimental Design	51
6.3	Results	51
7	Limitations, Recommendations and Conclusions	61
7.1	Introduction	61
7.2	Limitations and Recommendations	61
7.2.1	Random Body Motion (RBM) Cancellation	61
7.2.2	Radar Camera Calibration	62
7.2.3	Signal Quality Assessment	62
7.2.4	Participant Tracking	63
7.2.5	Robust Heart Rate Estimation	63
7.3	Conclusions	63

List of Figures

1.1	Outline of the structure of this dissertation.	10
2.1	Electrocardiography (ECG) electrodes placed on very specific positions on the body record a potential difference across time, with the characteristic PQRST shape. . . .	13
2.2	The skin over the muscle whose activity is to be monitored with EMG must be prepared.	17
3.1	Frequency of a typical chirp transmitted by an FMCW radar and the frequency of the corresponding echo received by the radar	20
3.2	The data received by each of the receivers in an FMCW radar can be represented as a cube called the radar data cube.	22
3.3	We imagine a target within the field of view of the radar with some displacement \vec{R}_0 at time $t = 0$	23
3.4	Two Gaussian noise corrupted range bin signal samples.	28
3.5	A demonstration of the effects of DC offsets and corrective measures on the radar phase signal.	29
3.6	An illustration of a time-indexed frequency estimate of the frequency of vibration of a speaker.	31
4.1	The radar phase signal encodes the small motion of the chest due to heart beat and respiration.	34
4.2	Heart rate signals from PPG (ground truth) and radar	35
4.3	The absolute errors between radar and PPG heart rate signals were calculated as a measure of performance.	37
5.1	Our respiration rate monitoring experimental setup consists of the Texas Instruments (TI) IWR6843ISK mmWave FMCW radar and the DCA1000EVM data capture board.	39
5.2	Our respiration rate estimation setup consists of a stereo camera pair and a radar between the cameras.	40
5.3	A rectified stereo pair of images is passed through readBarcode() to obtain the 2D pixel coordinates of a point of interest (POI).	41
5.4	Radar and camera data recorded during one of the respiration rate experiments.	43
5.5	Frequency spectra of the radar and camera data recorded during one of the respiration rate experiments.	45
5.6	Presented here are the average absolute errors of our respiration rate estimates averaged over experiments for each participant.	46
5.7	Camera phase signal versus radar phase signal scatter plots illustrating the relationship between the two signals.	48
6.1	Each participant sits on top of a table that is placed 70 cm in front of the tripod while wearing shorts to expose the QR code to the cameras.	51
6.2	During each experiment, three sets of time-synchronised radar, video and EMG data were recorded.	53

6.3	The exponential relationship we observe between normalised EMG and the corresponding normalised radar phase from a single contraction-relaxation cycle is a typical characteristic of muscles.	54
6.4	MATLAB's <code>lsqnonlin()</code> method from the Optimisation Toolbox is used to solve a curve-fitting problem to find values for coefficients A and B in the exponential equation.	55
6.5	Using Equation 6.1 and the model parameters in Figure 6.4, we computed the normalised radar phase predicted by the model given some normalised EMG from our data.	57
6.6	Presented here is the normalised EMG data from the illustrative experiment and illustrative contraction-relaxation cycle plotted against the corresponding normalised radar phase.	58
6.7	The camera phase signal correlates highly with the radar phase signal and has the same characteristic relationship with EMG that the radar phase signal has with EMG.	60
7.1	Micro-Doppler signature of speaker vibrating at 60 Hz with random body motion (RBM).	62
7.2	Micro-Doppler signature of the speaker after RBM cancellation.	63
7.3	Random Body Motion Cancellation Setup	64

Chapter 1

Introduction

In this dissertation, we demonstrate how a commercial off-the-shelf frequency modulated continuous wave (FMCW) radar and cameras can be used independently but complementarily, to monitor respiration rate, heart rate and muscle activity, cheaply and non-invasively. We do this by studying the subtle changes in the position (what we also call small motions) of appropriate parts of the skin across time in a radar return or video sequence.

1.1 Background

Some biological processes in the human body have the effect of inducing small amplitude motions of the surface of the skin. Such processes include: respiration, which causes periodic motion of the chest as air flows into and out of the lungs; the beating of the heart, which causes periodic motion of the skin (as well as the chest and head [1]) above arteries as blood flows through them; and muscle activity, which causes motion of the skin above the active muscle. These motions contain information about the processes that cause them.

This information is vital for assessing the state of health of an individual (hence the term vital signs). Traditionally, this information is obtained from data captured by sensors placed on or attached at an appropriate position on the skin of the person of interest. For respiration information, chest straps equipped with force sensors are often worn around the chest, while for heart beat a minimum of ten electrodes are placed on the subject's skin over the chest, arms and legs, in a technique called electrocardiography (ECG). Similarly for muscle activity information, electrodes are placed over the muscle of interest, in a technique called electromyography (EMG).

Though ECG and EMG are the current state of the art, they are expensive, invasive, uncomfortable, time-consuming, require an understanding of anatomy, training and the preparation of the subject's skin through shaving and cleaning with alcohol. Safe, easy to use, low-cost and non-invasive techniques for monitoring respiration, heart beat and muscle activity are needed.

1.2 Objectives and Scope

The objective of this research is to create a system to track biological processes in the human body by observing the minute movements they cause on the skin surface and underlying tissues. The created system should be able to monitor any biological process as long as it produces motion of sufficient magnitude and structure on the skin surface and underlying tissues. In comparison to current technologies for monitoring the respective biological processes, the system must be low-cost and operate without contacting the body. To demonstrate its performance, the developed system must be tested on some biological processes and compared to similar works in the literature.

One of the challenges that any non-contact method of monitoring small motion faces is the corruption of the small motion signal by motions with a larger magnitude. For small motions on the skin surface, the corrupting motions are usually any movements of the body, such as walking, the movement of limbs, leaning, turning, etc. Correcting for these random body movements is outside the scope of this work, though suggestions of how this can be addressed in future work are provided. Therefore, we make the assumption that the subject whose biological process we are interested in monitoring stays still throughout the monitoring process.

1.3 Outcomes

Owing to the remote sensing nature of radar and cameras, these were chosen as the sensing modalities for this project. Specifically, Texas Instruments' IWR6843ISK FMCW mmWave radar was used due to its ability to measure displacements with micron-level accuracy. Two GoPro HERO 7 Black cameras were used. Almost any pair of cameras could have been used. GoPros were used only because they were readily available in the laboratory. Because of the decision to use both radar and cameras independently, the developed system has the complementary strengths of the two sensors.

Cameras are readily available, easy to use and inexpensive. However, cameras suffer from occlusions e.g. by clothes, rely on good lighting conditions and the recording of people in video may infringe on their privacy. Though radar is more complicated to use, it can penetrate through clothing, performs well under poor lighting conditions and raises no privacy concerns. The radar component of the system costs \$640 while the camera component costs \$732.

The system's radar component transmits radio waves that are modulated by small motions on the skin and underlying tissues before being received by the radar. Analysis of this received signal reveals the information of interest. Specifically, a Fourier transform is applied on the first received chirp. From this, an estimate of how far the person is from the radar is made. This range estimate is then used to select and extract the phase of the correct Fourier coefficient. Extracting this phase from all the received chirps gives the phase signal which encodes the small motion. The camera subsystem uses triangulation to track the 3D position of a point on the skin. This position is then used to estimate the phase signal.

The performance of the developed system was tested on three applications: heart rate estimation, respiration rate estimation and muscle activity monitoring. In all our experiments, the participant was seated about 70 cm from the system. For our heart rate estimation experiments, a photoplethysmography (PPG) sensor was used to record ground truth data against which our radar system achieved 98.5% accuracy. For respiration rate estimation the participants were asked to breathe at certain frequencies, following an LED blinking at the given frequency. The radar system achieved 99.4% accuracy while the camera system achieved 99.3% accuracy. The Noraxon EMG system was used to provide ground truth for the muscle activity monitoring experiments. For this application, it was found that the extracted phase signal has the same characteristics as the muscle activity signal recorded through SMG, with the added benefit that ours is recorded without contact.

1.4 Overview

The structure of this dissertation can be thought of as comprising three parts: a preamble, experiments and results, and discussions. Figure 1.1 below illustrates the outline of this structure and the chapters that make up each part. The remainder of this section provides an overview of what each of the chapters in this dissertation entails.

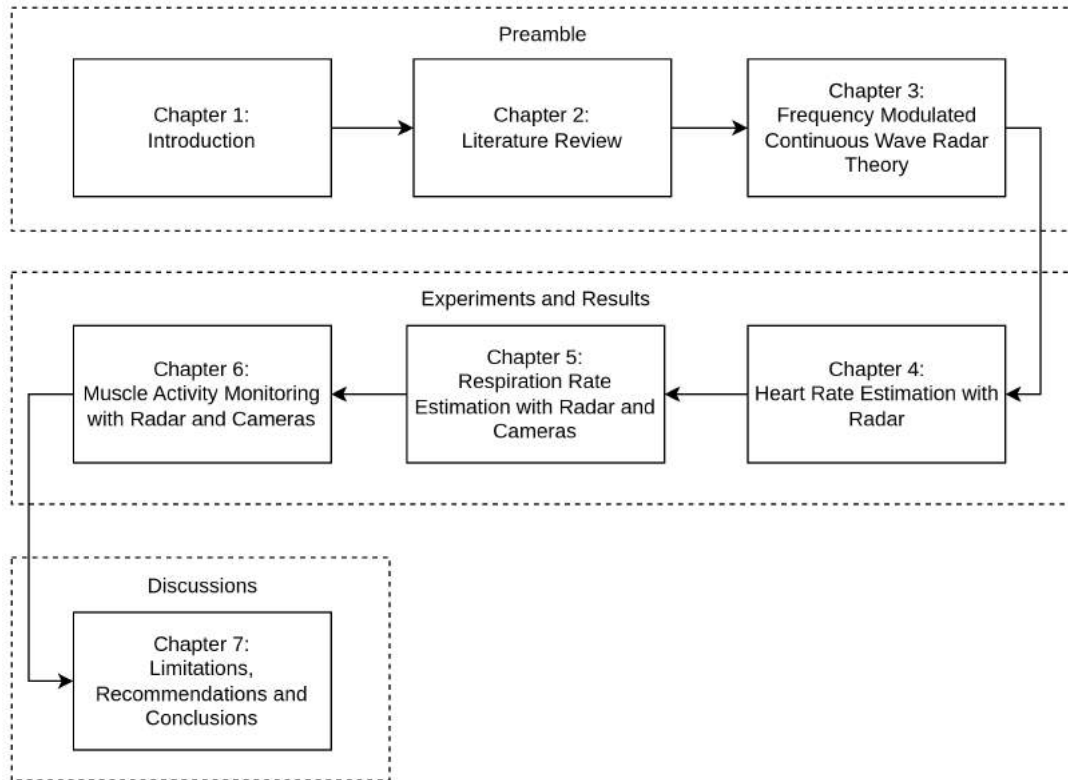


Figure 1.1: Outline of the structure of this dissertation.

1.4.1 Literature Review

In Chapter 2, we define each of the three physiological parameters we have chosen to monitor with our system. These parameters are heart rate, respiration rate and muscle activity. We give a brief description of how the underlying physiological or biological processes work. Additionally, we discuss the gold-standard or the current state-of-the-art way of monitoring that parameter. Finally, we explain how the parameter can be monitored without contact and discuss currently available approaches for doing so.

1.4.2 Frequency Modulated Continuous Wave (FMCW) Radar Theory

The basics of Frequency Modulated Continuous Wave (FMCW) millimetre wave (mmWave) radar are introduced in Chapter 3. A brief description of the characteristic format of FMCW radar data is then given. We then model a typical target and the small motion it exhibits. What follows is a model of the basic functional unit of FMCW radar – the chirp – based on the target’s small motion. We then discuss how this functional unit is used to estimate the target’s range using the Fourier transform. Finally, a discussion of how the target’s range is used to extract the small motion from the data is presented.

1.4.3 Heart Rate Estimation with Radar

Chapter 4 begins with an introduction to heart rate estimation with radar. This is followed by a description of the experimental design. That is the number of participants, description of the participants, the number of experiments conducted for each participant, the experimental protocol, the reference sensor used and how the sensors are time synchronised. Finally, the results found from analysing the data are presented and discussed.

1.4.4 Respiration Rate Estimation with Radar and Cameras

This chapter introduces respiration rate estimation with radar and cameras. What differentiates this chapter from the previous chapter technically, is that in this chapter, cameras are also used to monitor the small motion. Therefore, in addition to a description of the experimental design and the sensors used, we also give an in-depth description of how two cameras are used to mimic the radar. Finally, the results for both the radar and camera respiration rate estimation are given.

1.4.5 Muscle Activity Monitoring with Radar and Cameras

In this chapter, the need for a non-contact muscle activity monitoring approach is discussed. A brief introduction to the idea of using the muscle deformation signal as a measure of muscle activity is introduced. We then describe the experimental design and setup used to collect data. The results obtained from the data analysis are then presented and discussed, with comparisons made between our muscle activity signal and the signal acquired through SMG.

1.4.6 Limitations, Recommendations and Conclusions

In Chapter 7, we highlight important limitations of our non-contact small-motion monitoring approach. A potential way of addressing each limitation in future work is provided. We also draw an overarching picture of what this dissertation achieved. Special attention is paid to the significant parts of the work.

Chapter 2

Literature Review

2.1 Introduction

The beating of the heart, respiration, and contraction and relaxation of muscles are all important for the day to day functioning of our bodies. All three of these processes also induce characteristic and subtle motions on some areas of the skin. These processes therefore lend themselves very well as candidate processes to be monitored by the non-contact small-motion monitoring system presented in this dissertation. The three subsections below present the biological origin of these subtle motions as well as some approaches currently used to monitor these processes.

2.2 Heart Rate

All living cells require, amongst other things, oxygen, glucose and amino acids and also a way to excrete wastes such as carbon dioxide. For unicellular organisms this is trivial since the single cell is in direct contact with its environment and can achieve this through diffusion. In contrast, the vast majority of cells in multicellular organisms (such as humans) are not in direct contact with the environment, necessitating a system to deliver resources to and waste away from cells. The heart, blood and blood vessels make up this system, which is called the circulatory system [2].

The heart serves as the pump for the circulatory system. It receives deoxygenated blood from cells before pumping it to: the lungs to replenish oxygen and excrete carbon dioxide; the intestines to replenish nutrients; the liver to remove toxins from the blood; the kidneys to excrete wastes via urine and then finally to all the cells of the body [2].

To achieve its pumping action, the heart relaxes and contracts, resulting in the reception and ejection of blood, respectively [2]. Similarly to the muscle cells of skeletal muscles (see Section 2.4), cardiac muscle cells contract in response to action potentials. But unlike skeletal muscles, these action potentials do not originate from motor neurons but originate spontaneously from within the heart, specifically from the sinoatrial node or the pacemaker of the heart [2] [3].

The pumping or “beating” of the heart in this manner is rhythmic and is an intrinsic property of the heart [3]. The frequency of this rhythmic beating, usually expressed in number of beats in a minute (or bpm in short), is what we mean by heart rate. A normal heart rate for a resting human adult ranges from 60 bpm to 100 bpm, though it can be lower than 60 bpm for athletes and greater than 100 bpm during exercise [2].

Cardiac output is defined as the volume of blood the heart pumps out per unit time [3]. Sometimes the heart is unable to supply sufficient blood to the cells of the body, resulting in a condition called heart failure [2] [3]. Because cardiac output is a function of heart rate [2] [3], it is very im-

portant to monitor heart rate. Other life threatening conditions due to an increase or reduction in cardiac output (and thus may be signalled by heart rate) are hypertension and hypotension, respectively [2]. For these reasons, heart rate is an important signal for state of health and is referred to as a vital sign.

2.2.1 Electrocardiography (ECG)

As action potentials propagate through cardiac muscle cells, they cause electrical currents to spread throughout the body since tissues near the heart can conduct currents originating in the heart [2]. At any given time, the electrical activity originating in the sinoatrial node and propagating through the heart can be represented as a vector in three-dimensional space [2] [3]. A pair of electrodes placed at appropriate positions on the skin can measure the projection of this vector along the line joining the two electrodes [2] [3]. This technique is known as electrocardiography (ECG) and is used to monitor multiple factors of the heart, one of which is the heart rate.

ECG electrodes are placed on the arms, legs and on very specific positions on and around the left side of the chest (see Figure 2.1a). The limb electrodes measure projections in the frontal plane of the body while the chest electrodes measure projections in the transverse or horizontal plane, which is perpendicular to the frontal plane [2]. Each of the electrodes records a potential difference across time, with the characteristic shape in Figure 2.1b, reflecting the magnitude and direction of the underlying vector and the position of the electrode on the body's surface.

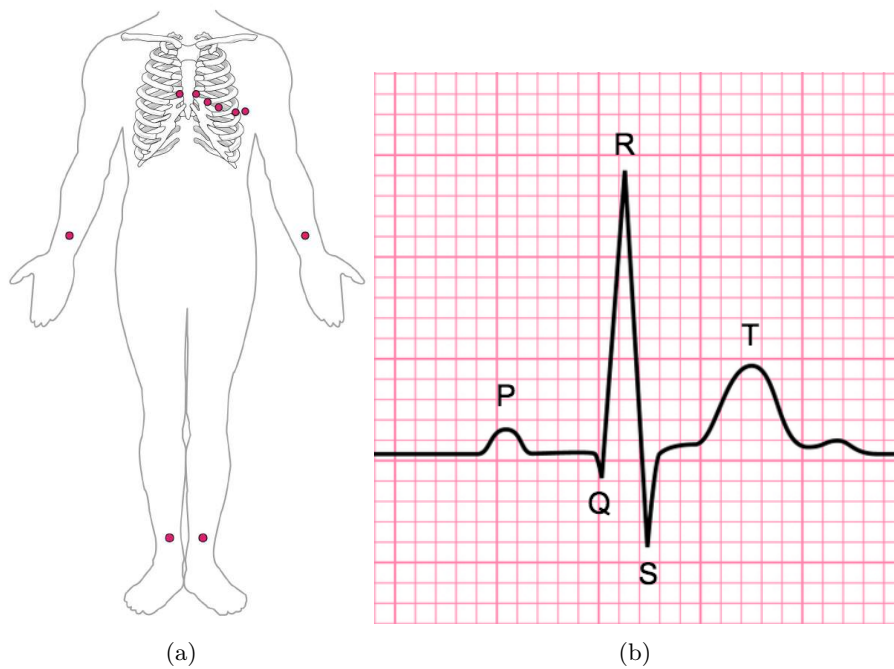


Figure 2.1: (a) Electrocardiography (ECG) electrodes are placed on the arms, legs and on very specific positions on and around the left side of the chest [4]. (b) Each of the electrodes records a potential difference across time, with the characteristic shape presented [5].

ECG is the current state-of-the-art in monitoring heart rate. However, as discussed above, it is a contact-based approach, requiring the subject to remove clothing from the torso and limbs. Additionally, if necessary, the subject may need to have the hair on their chest and limbs shaved for the electrodes to stick to the skin. All these factors render ECG an uncomfortable procedure. Furthermore, the electrodes on the chest must be placed at very specific positions, requiring a good

understanding of anatomy. This makes it impractical to monitor the heart rate for the elderly at home, those in palliative care and outpatients of one form or another.

2.2.2 Non-contact monitoring of heart rate

There are myriad non-contact ways of monitoring the heart and obtaining information, such as heart rate, about it. However, all these approaches work by measuring either motion or colour variations. The motion approaches measure either the minute periodic motion of the skin above arteries as blood flows through them [6], the displacement of the chest region overlying the apex of the heart [7] or the whole-body Newtonian reaction to the cardiac ejection of blood at each heart beat [1]. On the other hand, the colour variation approaches measure the subtle skin colour changes due to blood flowing into and out of blood vessels [8] [9].

As the heart beats, the chest displaces with amplitudes of about 400 microns while the skin on the leg displaces about 80 microns [6]. This displacement of the chest due to heart beat is about one to two orders of magnitude less than the displacement of the chest due to breathing (see Section 2.3). Additionally, a sensor's ability to record this signal is very sensitive to where the sensor is placed relative to the apex of the heart [7]. These factors make it challenging to both record and process the signal.

Techniques relying on motion use radio frequency based active sensors such as radar or WiFi [6]. These sensors transmit a wireless signal and recover the heart rate by comparing the transmitted signal to a version of this signal reflected off the subject whose heart rate is of interest. Some WiFi-based techniques must maintain a direct line-of-sight to the subject [10] whereas radar can penetrate through many types of materials including clothing [6]. The mmWave radar used by Zhao et al. is particularly cheap (less than \$100) and is capable of measuring displacements at micron-level accuracy [6].

The colour variation based techniques use RGB cameras to detect the subtle changes in skin colour due to blood flow [8] [9]. These techniques fall under remote photoplethysmography (rPPG) or photoplethysmography imaging (PPGi). Limitations of these techniques include: 1) the camera must always maintain direct line-of-sight with the subject, 2) camera performance is heavily dependent on lighting conditions (state-of-the-art rPPG techniques still make assumptions about the available light source [9]), 3) the use of cameras raises privacy concerns [6], and 4) these techniques still perform poorly on people with dark skin colour [9] [11] and on females [11].

2.3 Respiration Rate

In Section 2.2 above, it is mentioned that the heart pumps blood to the lungs to replenish it with oxygen and expel carbon dioxide. Respiration is the process through which the body takes oxygenated air from the environment into the lungs and ejects deoxygenated air out of the lungs. This process occurs in two stages, inhalation (or inspiration) and exhalation (or expiration). During inhalation, the lungs fill with oxygenated air through the nose and/or mouth while during exhalation, the lungs expel deoxygenated air through the nose and/or mouth.

Like the beating of the heart, respiration is rhythmic. We can therefore talk about a respiratory cycle, from the start of one inhalation to the end of the following exhalation. In normal resting human adults, this cycle usually takes about 3.75 s to 5 s, which translates to a respiration or breathing rate between 0.2 Hz and 0.27 Hz or 12 bpm and 16 bpm (adopting the same number of beats per minute convention as for heart rate). It is interesting to note that exhalation is slightly longer than inhalation; for a 5 s cycle, inhalation takes 2 s while exhalation takes 3 s [12].

2.3.1 Other respiration rate monitoring approaches

Unlike heart rate monitoring, which has the ECG as a gold-standard approach, respiration rate monitoring does not have a dedicated sensor. Respiration rate can easily be estimated by nurses just by observing the chest and/or stomach rising and falling over a minute or half a minute. Other ways of measuring respiration rate include using chest straps, breath humidity sensors, Wi-Fi [13] and motion sensors.

UbiBreathe [13] measures breathing rate by relying on the observation that the strength of a WiFi signal received by a WiFi-enabled device varies when the device is placed on an individual's chest. In general, the received signal strength (RSS) varies if an individual is on the line-of-sight between the device and the access point, even if the individual is not in contact with the device. UbiBreathe achieves an error of less than 1 breath per minute.

2.3.2 Non-contact monitoring of respiration rate

Unlike the heart, the lungs do not have muscles that relax and contract to bring about inhalation and exhalation [12] [14]. Instead, the lungs inflate and deflate in response to pressure changes within the thoracic cavity. These pressure changes are brought about mainly by the contraction and relaxation of the diaphragm, which is the major muscle of respiration [12]. Other muscles contributing to respiration are the external intercostal muscles, which pull the ribs upwards and forwards during inhalation [12]. As a result, the chest and lungs always move together during respiration [12]. Synchronicity between the diaphragm EMG (see Section 2.4) and chest wall movement due to respiration have been observed in a mouse [15].

This physiological connection between chest movement and respiration in time is the reason we can estimate breathing rate by monitoring the displacement of the chest. Importantly, the existence of sensors that use electromagnetic radiation to perform non-contact motion monitoring makes it possible to monitor respiration rate without contact.

Vital-Radio [16] uses a radar technique called FMCW (see Chapter 3) to measure the minute chest movements due to inhalation and exhalation. During inhalation, the volume of the chest cavity increases, the lungs expand, and the chest moves outwards away from the body. Conversely, during exhalation, the chest moves inwards, towards the body. Because of this motion, the distance between the chest and a wireless signal transmitter and receiver at a fixed position in front of the chest varies during breathing. By tracking how this distance varies with time, Vital-Radio estimates an individual's breathing rate. Vital-Radio tracks the phase of the received wireless signal, and since this phase is linearly related to the distance to the reflecting chest, Vital-Radio can estimate breathing rate.

Another non-contact respiration rate monitoring approach is marker-based motion capture. In addition to estimating respiration rate, this approach is often used to analyse chest wall kinematics, and breathing volume changes and to build chest wall models. This approach is particularly advantageous when the respiration rate is to be estimated in the presence of non-breathing movements. The typical setup consists of at least two, but usually more, cameras recording video of the subject. Between 32 and 89 hemispherical IR photo-reflective markers are placed on the subject's chest. The computation of the 3D trajectories of the markers across the video is expensive, which discourages the use of this approach [17].

All non-contact respiration rate sensors rely on being able to measure the movement of the chest due to inhalation and exhalation. In response to inhalation and exhalation, the chest experiences displacements in the range of 3 – 12 mm [18] [19]. Additionally, the displacement signal has been observed to have a shape resembling the voltage across a charging and discharging capacitor [18] [20]. This is because the rate at which the lungs fill with air during inhalation is highest at the beginning and drops as the lungs get full. Similarly, the lungs empty up much faster at the beginning than

at the end of exhalation. This characteristic non-sinusoidal shape introduces harmonics in the frequency spectrum of the chest displacement signal .

Finally, respiration is both an involuntary and a voluntary process. This means that we can, as we often do, breathe without voluntarily choosing to. However, we can also choose to breathe or not to. In Chapter 5 we take advantage of the voluntary nature of respiration to provide the reference signal we need to evaluate the performance of our respiration rate estimation approach.

2.4 Muscle Activity

Each skeletal muscle consists of a bunch of muscle fibers or muscle cells. Each of these muscle fibers in turn belongs to a muscle unit, which is a grouping of muscle fibers that are all innervated or excited by the same motor neuron [21].

During voluntary skeletal muscle contraction, the message instructing the muscle to contract propagates from the brain along motor neurons to the muscle. This message propagates as a difference in electric potential, known as an action potential, across the cell membranes of the motor neurons in its pathway. The motor neuron then discharges this action potential onto the muscle fibers it innervates, engaging their contractile proteins [21]. This process is what we mean by muscle activity.

Neurological and orthopaedic rehabilitation involves the treatment of patients suffering from disorders of the nervous system, bones and muscles. These disorders include spinal cord injury, stroke and multiple sclerosis. However, such rehabilitation is still out of reach for most South Africans, as the required equipment is prohibitively expensive. As a result, it is currently available only at biomechanics research laboratories and specialised centres. In South Africa, this is compounded by poor access to reliable transportation. This is unfortunate in a country where, as an example, the incidence of traumatic spinal cord injury (TSCI) is three times the global rate, with 60% of cases attributed to assault [22].

During rehabilitation, clinicians often need to know objectively if a patient's muscles are contracting. An electromyography (EMG) system provides clinicians with this knowledge. It achieves this by opening a window into the activities of muscles, such as contraction. Because of this, EMG is critical for rehabilitation from TSCI [23]. Electrical signals from the brain travel through motor neurons and cause the muscle fibers innervated by these neurons to be active [21]. EMG thus monitors muscle activity by measuring these electrical signals through surface electrodes placed on the skin over the muscle of interest (see Figure 2.2) or through needle electrodes inserted into the muscle.

2.4.1 Electromyography (EMG)

Action potentials are the source of all skeletal muscle activity and are also the source of the signals recovered through EMG. In this sense, EMG measures the cause of muscle activity. EMG measures muscle activity through a pair of conductive gel electrodes placed on the surface of the skin above the muscle of interest, as depicted in Figure 2.2b.

The EMG signal is affected by the tissues between the muscle and the electrodes. Additionally, the electrode-skin interface is intrinsically noisy and has an impedance. Ideally, these parameters should be minimised for the best signal quality. Larger surface area electrodes reduce both noise and impedance at the interface but introduce undesirable spatial low-pass filtering effects [21].

Prior to placing the electrodes, the skin is prepared to reduce the impedance and noise at the electrode-skin interface. Skin preparation involves shaving the skin to remove hair, as presented in




Figure 2.2: (a) Patch of skin over the gastrocnemius muscle shaved and cleaned with alcohol, and (b) EMG electrodes placed over the gastrocnemius muscle.

Figure 2.2a. Thereafter, the skin is rubbed with ethyl alcohol or abrasive conductive paste, washed with water and soap or stripped with adhesive tape to remove oily substances and the most superficial portion of the epidermis. Rubbing with abrasive conductive paste and rinsing is the most effective known treatment for reducing both the noise and impedance at the electrode-skin interface [21].

EMG requires contact with the subject's skin and is thus an invasive and potentially uncomfortable method for the measurement of muscle activity. Because of tissues between the muscle and electrodes, the EMG signal may be contaminated by muscle activity from other muscles near the muscle of interest. This is called crosstalk [21]. Moreover, the setup procedure has been termed time-consuming by EMG experts and is listed as one of the main barriers to the clinical employment of the technique [24].

2.4.2 Non-contact monitoring of muscle activity

When a muscle contracts, either voluntarily or during an electrically elicited contraction, the muscle vibrates and produces sound [25]. All non-contact approaches to monitoring muscle activity measure either the sound or the vibration produced by the contracting muscle. To differentiate it from EMG, this approach is called mechanomyography (MMG). The sound measured through MMG begins after the muscle action potentials [25]. In this sense, MMG measures the effects of muscle activity. MMG is to EMG what the non-contact approaches to measuring heart rate are to ECG. Indeed, many authors have suggested that MMG can be used as an alternative to EMG [26] [27] [28] [29].

Sensors that measure acceleration or displacement at the skin surface are used for MMG measurements. Examples of such sensors include accelerometers and piezoelectric sensors. Microphones are also used to measure muscle sound. It is important to note however that there is no single sensor that can be called an MMG sensor [30]. To ensure that the recorded MMG signal is not aliased, it is important to consider the sampling rate of the sensor used. The MMG signal is known to have negligible discernible energy above 50 – 60 Hz [31]. At a minimum, the sensor should sample at 120 Hz to satisfy the Nyquist-Shannon sampling criterion.

These mechanical sensors are attached to the skin. Relative motion between the sensor and the skin introduces artefacts in the recorded signal. Therefore, enough force should be applied to ensure secure attachment. At the same time, too much skin-sensor contact pressure or too heavy a sensor

will influence the dynamics of the muscle and affect the recorded signal [31] [25]. Like EMG, the MMG signal is affected by the tissue between the muscle and the sensor. The signal detected on the skin surface depends on the force applied by the muscle on the tissue and on the mechanical properties of the tissue [25]. To reduce noise and other artifacts, skin preparation may be necessary before sensor placement. Unlike EMG however, MMG is not affected by skin impedance and thus the required skin preparation may not be as intensive [32].

Rather than passively listening to the sound produced by muscle activity, ultrasound imaging has also been extensively used to actively monitor muscle activity by studying the morphological or architectural changes of muscles during contraction [33] [34] [35] [36] [37]. This approach is called sonomyography (SMG). When a muscle contracts, its sarcomeres (the basic contractile units of muscles) necessarily change length [21] [37]. Because the volume of muscles remains constant [38], it follows that muscle activity leads to shape change or deformation of the muscle [37]. It has been demonstrated that this deformation correlates with muscle activity [33] [34] [35] [37]. Most SMG algorithms measure this muscle deformation by measuring either muscle cross-sectional area, thickness [37] or muscle surface position [33] across a sequence of ultrasound images of the muscle.

SMG offers good spatial resolution, which solves the problem of muscle crosstalk [33] and can measure activity from deep seated muscles [35]. However, most SMG algorithms make use of computationally expensive tools such as optical flow to recover the SMG signals [34] [37]. Finally, SMG would be technically impractical to implement without contact because of the large difference between the acoustic impedance of air ($0.4 \times 10^3 \text{ kg}\cdot\text{m}^{-2}\cdot\text{s}^{-1}$ [39]) and that of human skin ($1.99 \times 10^6 \text{ kg}\cdot\text{m}^{-2}\cdot\text{s}^{-1}$ [40]), which would lead to most of the acoustic energy being reflected at the air-skin interface.

Evidently, contact-based MMG suffers most of the disadvantages of EMG. Laser Doppler Myography (LDMi) is the only non-contact MMG approach that, to the author's best knowledge, appears in the literature [41] [42]. LDMi uses the highly sensitive technique called Laser Doppler Vibrometry (LDV), commonly used to measure vibration, to measure the vibrations that accompany muscle activity. This technique measures the Doppler shift in frequency between an incident laser beam and a reference laser beam.

In contrast to microphones, accelerometers and piezoelectric sensors, the lasers used in LDMi do not mass load the skin and have no sensor resonance problems [41]. However, Laser Doppler Vibrometers are prohibitively expensive, costing even more than EMG systems [41]. Finally, the laser beam makes a point less than 1 cm^2 in area on the skin [41] [42]. This point must be kept at the same desired area, which greatly limits the application of LDMi to isometric contractions where the muscle contraction is not accompanied by gross motion [41].

2.4.3 Summary

ECG and EMG are the gold-standard multi-electrode approaches for monitoring heart rate and muscle activity, respectively. Though there is no one standard approach for monitoring respiration rate, there are a lot of contact and non-contact approaches for monitoring it. Some limitations of current camera-based respiration rate monitoring approaches are the use of many cameras (usually four), many markers (32 to 89) and complicated motion capture algorithms. There is no non-contact muscle activity monitoring approach that measures the characteristic muscle deformation signal seen in the SMG literature. In this dissertation, we seek to particularly address these gaps.

Chapter 3

Frequency Modulated Continuous Wave (FMCW) Radar Theory

3.1 Introduction

Frequency Modulated Continuous Wave (FMCW) is a specific type of modulation scheme used for active sensors. It is often implemented using radio waves with wavelengths in the millimetre range, to form an mmWave FMCW radar. The relatively short wavelengths allow mmWave radars to be compact and low-cost [19]. Additionally, the frequency modulation allows for wide bandwidths and thus fine range resolution. Finally, the FMCW design opens up a host of tunable parameters, making the radar flexible to different measurement scenarios [19].

As the name suggests, the frequency of the transmitted wave in an FMCW radar is modulated. A sawtooth modulation scheme is often used [19]. This means that the frequency of the transmitted signal, $f_T(t)$, varies linearly with time according to:

$$f_T(t) = \frac{B}{T}t + f_0, \quad (3.1)$$

where B is the bandwidth of the frequency sweep, T is the time it takes to sweep the bandwidth and f_0 is the starting frequency.

This signal with a linearly increasing frequency is often called a chirp. After the sweep time, T , the frequency starts from the start frequency, f_0 , again. The frequency of the transmitted signal repeats this as many times as the required number of chirps. For the same sweep time, T , increasing the number of chirps increases the observation time and thus leads to a finer velocity resolution.

Unlike conventional radars, FMCW radars do not estimate the range of a target by directly measuring the time it takes the transmitted signal to reflect off the target and be received back [19]. Instead, they estimate range by calculating the difference between the frequency of a transmitted chirp and the frequency of this chirp's echo, reflected off the target. Imagine a chirp with the frequency in Equation 3.1 is transmitted, it reflects off a target at a distance R from the radar and it is received by the radar after time t_d . Because the chirp travels a distance R to the target and the echo travels the same distance back, time t_d is equal to

$$t_d = \frac{2R}{c}, \quad (3.2)$$

where c is the speed of electromagnetic radiation.

The frequency of the received echo, $f_R(t)$, will be equal to the frequency of the transmitted chirp but delayed by time t_d . Mathematically,

$$f_R(t) = f_T(t - t_d) = \frac{B}{T}(t - t_d) + f_0. \quad (3.3)$$

Graphically, the frequencies of the transmitted chirp and the received echo can be represented as illustrated in Figure 3.1 below.

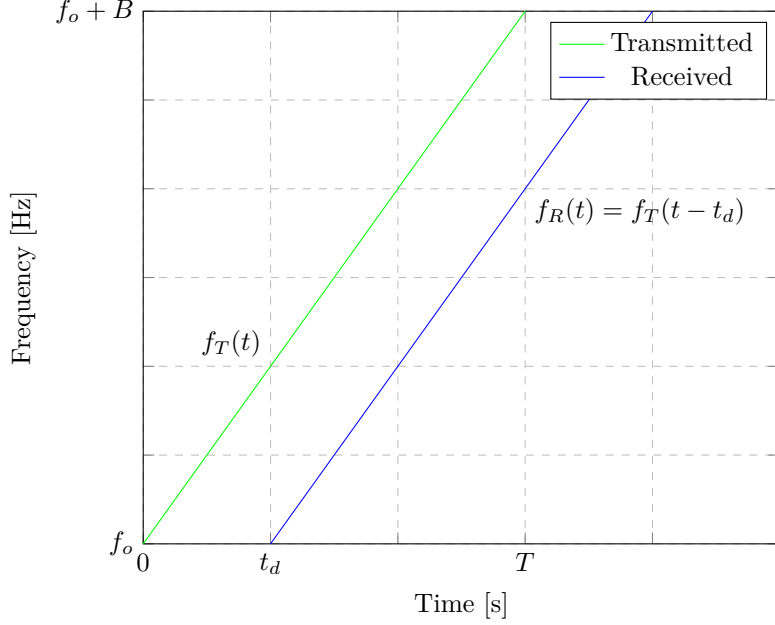


Figure 3.1: Frequency of a typical chirp transmitted by an FMCW radar and the frequency of the corresponding echo received by the radar time t_d after the chirp reflected off a target a distance R from the radar.

The difference between the two frequencies in Figure 3.1 is called the beat frequency, $f_b(t)$. For our scenario, the beat frequency is

$$f_b(t) = f_T(t) - f_R(t) = \frac{B}{T}t + f_0 - \left[\frac{B}{T}(t - t_d) + f_0\right] = \frac{B}{T}t_d = \frac{B2R}{Tc}. \quad (3.4)$$

The signal at the output of the radar's ADC is at this frequency (refer to Equation 3.12 modelling a chirp received by an FMCW radar). Once the beat frequency is known, Equation 3.4 can be used to calculate the distance, R , to the target. Moreover, the maximum detectable range, R_{max} , can be calculated using Equation 3.4, given the maximum sampling frequency of the radar's ADC. In fact,

$$R_{max} = \frac{cT}{2B} \frac{f_{smax}}{2}, \quad (3.5)$$

where f_{smax} is the maximum sampling frequency of the radar's ADC ¹.

In practice, the Discrete Fourier Transform (DFT) is used to estimate the range of a target from FMCW radar data. Of course, a variation of the computationally efficient Fast Fourier Transform (FFT) is actually used. As mentioned already, multiple chirps are often transmitted. The FFT is then applied on each of the received echoes and the result of such a computation is called a Range FFT. It is called this because the resulting frequency axis can be interpreted as range and

¹Note that the second 2 in Equation 3.5 is there to respect the Nyquist-Shannon sampling criterion. In actuality, for example for Texas Instruments' FMCW radars, the maximum sampling frequency of the radar's ADC is further multiplied by 0.9 [43].

the magnitude of the Fourier coefficients are a measure of how much energy is reflected at that range. Importantly, the magnitude of the Fourier coefficient at the beat frequency will be maximal amongst the magnitudes of the Fourier coefficients (see Section 3.5 about range estimation). From this fact, we can determine where the target is using the beat frequency equation, Equation 3.4.

3.2 FMCW Radar Data

It is often useful to think of FMCW radar data as a multidimensional data structure (see Figure 3.2)². Each dimension allows us to measure a different physical parameter. We have already seen how a single chirp — the smallest functional unit of the data — can be used to measure range. This dimension of the data, the dimension indexed by the samples of a chirp, is called fast time (because of the relatively higher sampling rate of the ADC compared to other dimensions). The dimension indexed by chirps is called slow time and it allows us to measure velocity and thus Doppler frequency shift. Chirps are divided into groups called frames and the dimension indexed by frames is simply a time dimension, allowing the observation of how the other parameters vary with time.

Furthermore, in FMCW radars, which often feature multiple receivers and transmitters, the receivers are spaced apart on the circuit board. This spatial separation results in differences in the time it takes for the reflected radio waves to reach each receiver [44]. This time difference makes it possible to measure the angle of arrival of the radio waves. In this sense, the dimension indexed by the receivers allows us to measure angle.

Using the data from each of the receivers in an FMCW radar, a three-dimensional data structure as depicted in Figure 3.2 can be formed. This is referred to as a radar data cube. One radar data cube can be formed for each receiver.

3.3 Small Motion Model

In view of the fact that we are interested in monitoring small motions that may (heart rate and breathing rate) or may not (muscle activity) be periodic, here we model what that motion might look like. Let us consider a target that is initially, when the radar starts transmitting at time $t = 0$, a displacement \vec{R}_0 from the radar, sitting at some point A within the field of view (FOV) of the stationary radar (see Figure 3.3). For the IWR6843ISK FMCW radar considered here, the x and z axes correspond to the azimuth and elevation axes, respectively [45]. For our purposes, the target could be the chest, the thigh or any other appropriate surface (see Chapters 4, 5 and 6).

As the radar is recording data, the target is free to move within the FOV exhibiting the small motion. At any time $t \geq 0$, let the target's displacement from point A be \vec{X}_t . Then, the displacement of the target from the radar at time t , \vec{R}_t , becomes

$$\vec{R}_t = \vec{R}_0 + \vec{X}_t. \quad (3.6)$$

The magnitude of this displacement, $|\vec{R}_t|$, is the target's range that the radar measures as described in Section 3.1 above and Section 3.5 below. However, there is a limitation on estimating range this way. If the magnitude of the small motion, $|\vec{X}_t|$, is less than the range resolution, ΔR , then this small motion will not change the estimated range i.e. the small motion will not be detected by the radar. In such a case, we say the target moves within a range bin. To solve this problem, the beat frequency obtained from the Range FFT of the first chirp echo is used to estimate where the target roughly is, i.e. to estimate $|\vec{R}_0|$. Then the phase of the Range FFT coefficient at this beat frequency (or range) is calculated for all chirp echoes [46]. The sequence of all these phases is an

²Each building block in the radar data cube representation in Figure 3.2 is a sample belonging to a chirp. Notice that there are only 10 samples per chirp here. In practice, however, each chirp usually has a lot more samples.




Figure 3.2: The data received by each of the receivers in an FMCW radar can be represented as a cube called the radar data cube.

estimation of the small motion, specifically the component of \vec{X}_t along \vec{R}_0 i.e. along the radar's line of sight or the radial direction [19].

The velocity of the target at time t , \vec{V}_t , is

$$\vec{V}_t = \dot{\vec{R}}_t = \dot{\vec{R}}_0 + \dot{\vec{X}}_t = \dot{\vec{X}}_t. \quad (3.7)$$

However, continuous wave radars estimate the target's velocity by measuring the shift in the received echo's frequency as compared to the frequency of the transmitted chirp [19]. This shift is due to the Doppler effect and is caused by the component of the target's velocity that is radial to the radar [19]. With this approach, this implies that the radar would only be able to estimate the component of \vec{V}_t that is resolved along the y-axis in Figure 3.3.

Similar to range estimation, this radial velocity is estimated through an FFT operation called the Doppler FFT. In contrast to the Range FFT, the Doppler FFT is applied on the complex-valued output of the Range FFT, not on the complex-valued raw received echo data. Also unlike the Range FFT, the Doppler FFT is not applied along the fast time dimension but along the slow time dimension.

3.4 Received Echo Model

We are interested in using mmWave FMCW radar to monitor small motions due to certain biological processes at the surface of the skin. In the previous sections, we have discussed how a frequency

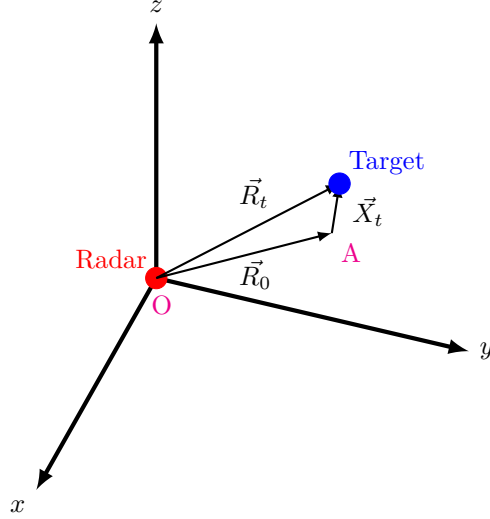


Figure 3.3: We imagine a target within the field of view of the radar with some displacement \vec{R}_0 at time $t = 0$. At any time $t \geq 0$, the target is free to have some small displacement \vec{X}_t from point A.

analysis of the received echo (or echoes)³ can reveal where the surface of the skin roughly is, how it moves subtly across time and how fast it is moving. Here, we briefly present a mathematical model of what the transmitted chirps and received echoes look like.

As seen in [46], a typical chirp transmitted by an FMCW radar can be represented as

$$S_T(t) = A_T e^{j(2\pi f_0 t + \pi K t^2) + \phi_0}, 0 \leq t \leq T, \quad (3.8)$$

where $K = \frac{B}{T}$ is the slope of the transmitted frequency (see Figure 3.1), ϕ_0 is the initial phase and A_T is the amplitude of the transmitted chirp.

This chirp is transmitted and reflects off the target (see Figure 3.3). Let $|\vec{R}_0| = R_0$ and let $x(t)$ be the component of \vec{X}_t along the y-axis. Then the received chirp, $S_R(t)$, after time t_d becomes

$$S_R(t) = S_T(t - t_d) = A_R e^{j(2\pi f_0 (t - t_d) + \pi K (t - t_d)^2) + \phi_0}, 0 \leq t \leq T, \quad (3.9)$$

where A_R is the amplitude of the received chirp and the time delay, t_d , is just as in Equation 3.2 but with the additional small motion $x(t)$:

$$t_d = \frac{2(R_0 + x(t))}{c}. \quad (3.10)$$

The amplitude of the received chirp, A_R , depends on multiple factors such as the peak transmit power of the transmitter, transmit antenna gain, receiver antenna gain, radar cross section (RCS) of the target, etc. [19].

In the receiver, the signal $S_R(t)$ goes through a number of amplifier stages, a down-conversion stage or stages (depending on whether the receiver has a homodyne or heterodyne architecture) and a low-pass filter [19] [46]. At the output is the baseband signal $S_B(t)$ in Equation 3.11 [46], given as an in-phase component (I) and a quadrature component (Q):

$$S_B(t) = S_T(t) \cdot S_R(t)^* = A_R e^{j(2\pi K t_d t + 2\pi f_0 t_d)}, 0 \leq t \leq T, \quad (3.11)$$

³We avoid using just the term chirp as this is often confusing when wanting to distinguish transmitted from received chirps. Instead, we use received chirp, received echo or echo for what is received by the radar and transmitted chirp for what is transmitted by the radar.

where $(\cdot)^*$ is the complex conjugate operator.

The term associated with t_d^2 has been ignored since $t_d^2 \ll t_{dt}$ [46]. Additionally, it has been assumed that during the chirp duration, T , the displacement of the target is negligible and thus the target has no velocity during this time [46]. If this assumption is not made (or is unreasonable), then the frequencies in both $S_R(t)$ and $S_B(t)$ should have a Doppler frequency shift term added to them [19] [46].

Now let us assume that the radar transmits and receives a total of M chirps, each with index $i \in \{1, 2, 3, \dots, M\}$. Also suppose that each received chirp has N samples. Then any received chirp can be generally described by

$$\begin{aligned} S_B(iT + t) &= A_R e^{j(\frac{4\pi K R_0}{c} t + \frac{4\pi f_0 R_0}{c} + \frac{4\pi f_c}{c} x(iT))} \\ &= A_R e^{j(2\pi f_b t + \frac{4\pi f_0 R_0}{c} + \frac{4\pi f_c}{c} x(iT))} \\ &= A_R e^{j(2\pi f_b t + \frac{4\pi R_0}{\lambda_0} + \frac{4\pi x(iT)}{\lambda_c})} \\ &= A_R e^{j(2\pi f_b t + \varphi_i)}, \end{aligned} \tag{3.12}$$

where $f_c = f_0 + \frac{B}{2}$ is the center frequency, λ_c is the center wavelength, $f_b = \frac{2BR_0}{cT}$ and $\varphi_i = \frac{4\pi R_0}{\lambda_0} + \frac{4\pi x(iT)}{\lambda_c}$.

It is thus clear that the phase of each received chirp, φ_i , encodes within it a sample of the small motion at time iT , $x(iT)$. Just as importantly, we see that the small motion is sampled at frequency $\frac{1}{T}$. We can think of this as the slow time frequency. Because of the Nyquist-Shannon sampling criterion, this frequency (or equivalently, period) places an upper limit on how fast the target can move while still being trackable with this approach (see Equation 3.24 for an expression of this maximum velocity).

3.5 Range Estimation

Our small motion model assumes that the target starts at some distance R_0 and then exhibits small motion $x(t)$. In the preceding sections, we have:

1. Mentioned that the Range FFT can be used to estimate R_0 (or equivalently, f_b).
2. Demonstrated that at any time t , $0 \leq t \leq T$, the phase of each received chirp with index i , $1 \leq i \leq M$, encodes the small motion we are interested in, $x(iT)$.

Here, we will demonstrate why 1 is true and then demonstrate that the phase of the Range FFT coefficients also encodes the small motion and we can use R_0 (or f_b) to find it.

As mentioned in Section 3.1, the Range FFT is simply the DFT applied on each received chirp, $S_B(iT + t)$. Mathematically, this is

$$X_k = \sum_{n=0}^{N-1} x_n e^{-j2\pi \frac{k}{N} n}, \tag{3.13}$$

where k is the frequency index, $x_n = S_B(iT + n \cdot T_s)$, $n \in \{0, 1, 2, \dots, N-1\}$ is the discretised received chirp and $T_s = \frac{1}{f_s}$ is the period at which the received chirp has been sampled, also referred to as the fast time sampling period.

Dividing and multiplying by T_s in Equation 3.13 gives

$$X_k = \sum_{n=0}^{N-1} x_n e^{-j2\pi \frac{k}{NT_s} n T_s}. \tag{3.14}$$

If we let $\frac{k}{NT_s} = \frac{k}{N}f_s = f_k$ and substitute into Equation 3.14 the correct expression for x_n , we get

$$\begin{aligned} X_{f_k} &= \sum_{n=0}^{N-1} A_R e^{j2\pi f_b n T_s} e^{-j2\pi f_k n T_s} e^{j\varphi_i} \\ &= A_R e^{j\varphi_i} \sum_{n=0}^{N-1} e^{j2\pi(f_b - f_k)n T_s}. \end{aligned} \quad (3.15)$$

If $f_k = f_b$, then

$$X_{f_k} = N A_R e^{j\varphi_i} \quad (3.16)$$

because $\sum_{n=0}^{N-1} e^{j2\pi(f_b - f_k)n T_s} = N$. Therefore, at the beat frequency, the Range FFT coefficient has magnitude $|X_{f_k}| = N A_R$.

However, if $f_k \neq f_b$, then we can let $z = e^{j2\pi(f_b - f_k)T_s}$, $z \in \mathbf{C}$. Then by noticing that our sum is a geometric sum, we have

$$\sum_{n=0}^{N-1} e^{j2\pi(f_b - f_k)n T_s} = \sum_{n=0}^{N-1} z^n = 1 + z + z^2 + z^3 + \dots + z^{N-1}. \quad (3.17)$$

Letting $|\sum_{n=0}^{N-1} e^{j2\pi(f_b - f_k)n T_s}| = N^*$ and applying the triangle inequality on Equation 3.17 gives

$$N^* = \left| \sum_{n=0}^{N-1} e^{j2\pi(f_b - f_k)n T_s} \right| \leq |1| + |z| + |z^2| + \dots + |z^{N-1}| = N. \quad (3.18)$$

Therefore, at any frequency that is not the beat frequency, the Range FFT coefficient has magnitude $|X_{f_k}| = N^* A_R$. But Equation 3.18 states that $N^* \leq N$. Therefore, the magnitude of the Range FFT coefficient at the beat frequency is maximal i.e. there will not be a Range FFT coefficient with magnitude greater than this one at any other frequency. This is what allows us to determine where the target roughly is i.e. R_0 . Often, in practice, $|X_{f_k}| = N A_R$ is not only maximal but is the maximum, which makes finding R_0 easier. If an element is maximal, say, in a set, then there is no other element in that set that is larger than it. Maximum on the other hand implies that it is the largest element. The maximum element is also maximal but the converse is not necessarily true.

Additionally, Equation 3.16 shows that the phase of the Range FFT coefficient at the beat frequency is φ_i , which encodes $x(iT)$. At any other frequency f_k that is not the beat frequency, the phase of the Range FFT coefficient will be φ_i corrupted by the phase of $\sum_{n=0}^{N-1} e^{j2\pi(f_b - f_k)n T_s}$. We call the sequence of all phase samples, φ_i , across all received chirps the phase signal, φ . Concretely, the phase signal is

$$\varphi = (\varphi_i)_{i=1}^M. \quad (3.19)$$

This is the signal we extract from the data and process in all our applications.

One last thing to consider with range estimation is the range resolution. This is how far apart two targets should be for our Range FFT range estimation approach to be able to recognise them as two separate targets. A very fine range resolution is not necessary for our applications since we

do not use range but phase to monitor the small motion. However, we are concerned about range resolution for two reasons.

Firstly, we want our range resolution to be coarse enough for our target to remain within the same range bin for all time t , $0 \leq t \leq M \cdot T$. This means that determining the beat frequency from the Range FFT of the first received chirp should be enough. It should not be necessary to keep updating the frequency at which to look for the phase for each received chirp. This should be easy to achieve since the finest range resolution the mmWave FMCW radar we use can achieve is 37.5 mm ⁴ and the motion of the chest due to say, breathing, is in the range $3 - 12 \text{ mm}$ [18].

Secondly, the range resolution should be fine enough for us to be able to easily confirm that there is only the target of interest in its range bin. This is so that the motion encoded in the extracted phase signal is only due to the target of interest. Again, this should not be a problem since we can achieve a range resolution of 37.5 mm at best.

Having done all the work with the Range FFT, arriving at the expression for the range resolution, ΔR , is easy. The frequency resolution of the Range FFT, Δf_R , determines the range resolution. This frequency resolution is the difference between any two consecutive frequencies. That is,

$$\Delta f_R = f_{k+1} - f_k = \frac{k+1}{NT_s} - \frac{k}{NT_s} = \frac{1}{NT_s}, \quad (3.20)$$

but $NT_s = T$, therefore

$$\Delta f_R = \frac{1}{T}. \quad (3.21)$$

Mapping from the frequency domain to the range domain using the expression for the beat frequency in Equation 3.4 gives

$$\Delta R = \Delta f_R \cdot \frac{cT}{2B} = \frac{1}{T} \cdot \frac{cT}{2B} = \frac{c}{2B}. \quad (3.22)$$

3.6 Phase Analysis

In our signal processing pipeline, before getting the phase signal, φ , we first extract the Range FFT coefficient at the beat frequency for all received chirps. We call the sequence of all these coefficients the range bin signal, s . From Equation 3.16, the range bin signal is

$$s = (NA_R e^{j\varphi_i})_{i=1}^M = (I_i + jQ_i)_{i=1}^M. \quad (3.23)$$

At this point, the phase signal can be obtained from the range bin signal. A variety of ways for doing this exist. One simple and common approach is arctangent demodulation [19], which is simply taking the arctangent of the imaginary part, Q_i , over the real part, I_i , for each sample of the range bin signal. Because the phase samples, φ_i , in the range bin signal are in the range $[-\pi, \pi]$ (we say the phase is wrapped), after arctangent demodulation it is often necessary to unwrap the phase. This is necessary because if say $x(iT) = \frac{\lambda_c}{4}$, then $\varphi_i = \frac{4\pi R_0}{\lambda_0} + \pi > \pi$ for $R_0 \neq 0$ (see Equation 3.12 modelling a chirp received by an FMCW radar). Phase unwrapping works by making an assumption about the underlying motion [47].

The assumption is that the target does not exceed the maximum permissible velocity, v_{max} (see Equation 3.24 for the expression of v_{max}). From Equation 3.12, recall that $\varphi_i = \frac{4\pi R_0}{\lambda_0} + \frac{4\pi x(iT)}{\lambda_c}$.

⁴Calculated using Equation 3.22 and a bandwidth of 4 GHz [43].

⁵ $NT_s = T$ assumes that the radar's ADC samples the received chirp for the entire chirp time T . This is not exactly true [43], but it suffices as an approximation.

Notice that if, $x((i+1)T) = x(iT) + \frac{\lambda_c}{2}$, then $\varphi_{i+1} = \frac{4\pi R_0}{\lambda_0} + \frac{4\pi x(iT)}{\lambda_c} + 2\pi = \varphi_i$. Even though the small motion has changed by $\frac{\lambda_c}{2}$ over a time duration of $2T$, the phases do not capture this motion. Therefore, we conclude that over a time equal to two chirp times, $2T$, the small motion should change less than $\frac{\lambda_c}{2}$. This argument actually gives us an expression for the maximum velocity, v_{max} , the target can have before it becomes impossible to encode its small motion in the phase. We have

$$v_{max} = \frac{\frac{\lambda_c}{2}}{2T} = \frac{\lambda_c}{4T} = \frac{\lambda_c}{4T}. \quad (3.24)$$

One interpretation of Equation 3.24 is that the largest displacement the target can have in a time equal to the chirp time, T , is $\frac{\lambda_c}{4}$. So, given $\varphi_i = \frac{4\pi R_0}{\lambda_0} + \frac{4\pi x(iT)}{\lambda_c}$ and $\varphi_{i+1} = \frac{4\pi R_0}{\lambda_0} + \frac{4\pi x((i+1)T)}{\lambda_c}$, we have $\Delta\varphi = \varphi_{i+1} - \varphi_i = \frac{4\pi R_0}{\lambda_0} + \frac{4\pi x((i+1)T)}{\lambda_c} - \frac{4\pi R_0}{\lambda_0} - \frac{4\pi x(iT)}{\lambda_c} = \frac{4\pi \Delta x}{\lambda_c}$. Therefore, if the largest Δx can be is $\frac{\lambda_c}{4}$, then the largest $\Delta\varphi$ can be is π . Because the target can move towards and away from the radar, we also consider that the displacement Δx cannot be less than $-\frac{\lambda_c}{4}$. This then suggests that the smallest $\Delta\varphi$ can be is $-\pi$. We have

$$-\pi \leq \varphi_{i+1} - \varphi_i \leq \pi. \quad (3.25)$$

This result leads to the phase unwrapping algorithm proposed in [47]. The algorithm works as follows: iterate over all the wrapped phases. If the difference between two consecutive phases satisfies Equation 3.25, do nothing. If instead $\varphi_{i+1} - \varphi_i > \pi$, then subtract 2π from φ_{i+1} . Finally, if $\varphi_{i+1} - \varphi_i < -\pi$, then add 2π to φ_{i+1} .

Though retrieving the phase signal by performing arctangent demodulation and then applying the phase unwrapping algorithm above are theoretically simple operations, in practice they are often challenging. For noisy samples of the range bin signal, the arctangent demodulation gives incorrect phase estimates. Having incorrect phase estimates at the arctangent demodulation stage presents a considerable challenge to the subsequent phase unwrapping stage and the resulting phase signal.

This is because the decisions made in the phase unwrapping algorithm (add 0, 2π or -2π) rely on these phase values being correct. If the wrapped phase sample at index $i = k$ is sufficiently incorrect, then incorrect decisions at both indices $i = k$ and $i = k + 1$ will be made during the unwrapping process. This will lead to spurious jumps in the unwrapped phase signal at these samples. Moreover, these wrong decisions will render the unwrapped phase samples at all the indices $i \geq k$ incorrect. If enough of the wrapped phase samples are incorrect, the unwrapped phase signal will drift with time [19] [48].

We found that we can detect the samples in the range bin signal that would lead to incorrect wrapped phase estimates. The magnitude of each range bin signal sample, $|I_i + jQ_i| = NA_R$, provides a reliable signal of whether or not that sample will lead to an incorrect wrapped phase estimate. If the magnitude is relatively small, then the corresponding phase estimate will be incorrect. This is the same conclusion reached in [49], though in a very similar but different context. Figure 3.4 below has been reproduced from [49] as it depicts very well why $|I_i + jQ_i|$ is a good measure of the confidence in the corresponding phase estimate.

Figure 3.4 presents two range bin signal samples, one with a large magnitude and the other with a small magnitude. Depicted around each is a cloud of the possible complex numbers that each could take on if it were corrupted by Gaussian noise. Notice that the small magnitude sample can have almost any phase in the range $[0, 2\pi]$ whereas the large magnitude sample's phase remains fairly limited around whatever the correct phase value is (approximately $\frac{\pi}{4}$).

The Differentiate and Cross Multiply (DACM) algorithm offers an alternative phase retrieval approach to the arctangent demodulation discussed here [50]. With DACM, phase unwrapping is

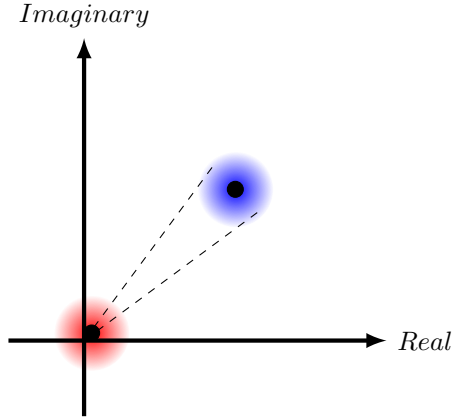


Figure 3.4: Two Gaussian noise corrupted range bin signal samples. Notice that the small magnitude sample can have virtually any phase in the range $[0, 2\pi]$, whereas the large magnitude sample's phase remains fairly limited around the correct phase value (approximately $\frac{\pi}{4}$).

unnecessary as the algorithm unwraps the phase automatically. DACM computes the unwrapped phase samples as follows [51]:

$$\varphi_i = \sum_{i=2}^M \frac{I_i[Q_i - Q_{i-1}] - Q_i[I_i - I_{i-1}]}{I_i^2 + Q_i^2}. \quad (3.26)$$

Even the DACM algorithm performs poorly when $|I_i + jQ_i|$ is small. Equation 3.26 presents $|I_i + jQ_i|^2$ as a divisor. If this number is very small, the quotient becomes very large. This results in a phase signal that can be orders of magnitude larger than it should be. Additionally, the range bin signal often needs to be interpolated and a 5-point differencing scheme is used instead of backward differencing [50] for DACM to perform as well as arctangent demodulation, thereby increasing computational complexity. Therefore, for our applications, arctangent demodulation was used.

Recall that at the receiver output, the complex-valued baseband signal, $S_B(t)$, in Equation 3.11 comprises a real and an imaginary part or an in-phase (I) and a quadrature (Q) component, respectively. Ideally, the I and Q components are sinusoids, with a phase offset of 90° and with equal amplitudes [19]. In reality, due to imperfections in the receiver's electronics, amplitude and phase imbalances exist between the two channels [19] [52] [53]. Additionally, the range bin signal often contains DC offsets due to coupling effects and environmental interferences [48]. Both of these negatively affect the ability to recover the phase signal and thus must be corrected for.

Various methods exist for IQ imbalance and DC offset correction [19] [47] [52] [53]. Through experimenting with the different algorithms, we found that the DC offset correction algorithms [19] [47] were very effective at improving phase signal retrieval whereas the IQ imbalance algorithms [19] [53] had no positive effect. Two DC offset correction algorithms were specifically investigated. One corrected the DC offset by subtracting the mean of the range bin signal from the range bin signal (see Figure 3.5c and d). Instead of the mean, the second one subtracted the centre of the range bin signal's scatter plot (see Figure 3.5e and f). This centre was estimated using a least squares approach [47]. Unless stated otherwise, all our radar data was recorded with chirp sweep time of $1670.17 \mu\text{s}$, 2048 ADC samples, 255 chirps, 40 frames, frame periodicity of 1296 ms and we make use of time-division multiplexing multiple-input-multiple-output (TDM-MIMO) (see Chapter 4). These settings result in a bandwidth of 2.45 GHz, which corresponds to a range resolution of 6.13 cm and a velocity resolution of 0.19 cm/s.

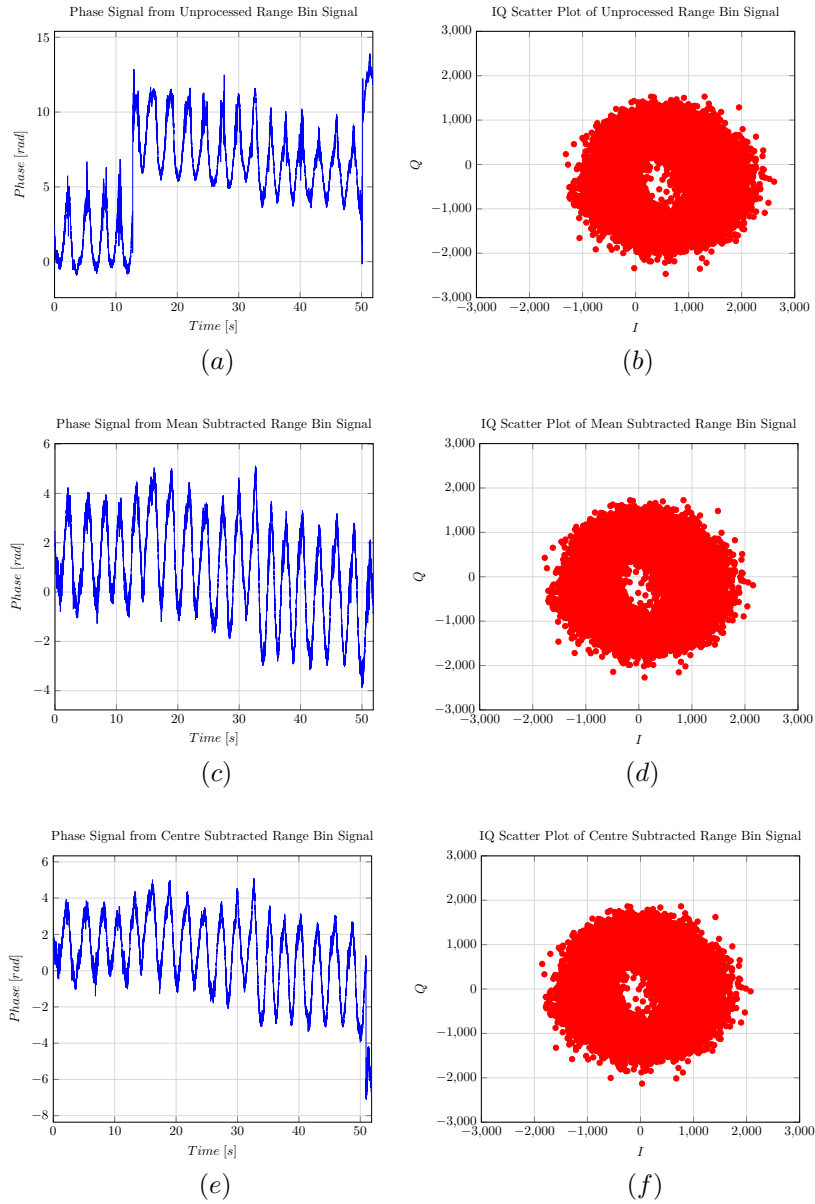


Figure 3.5: (a) Presents the radar phase signal extracted from the unprocessed range bin signal in (b). The phase signal is corrupted by artefacts because of DC offsets in the range bin signal. The artefact-free phase in (c) is obtained from (d), the range bin signal after mean subtraction. In (e) is the phase obtained from the centre subtracted range bin signal in (f).

Figure 3.5a delineates that the phase signal extracted from the unprocessed range bin signal is negatively affected by spurious jumps at about 12 s and at 50 s. Figure 3.5c illustrates that both jumps can be corrected by a simple IQ mean subtraction. The results of the more complicated centre subtraction approach in Figure 3.5 are good but there is a spurious jump at about 50 s. In this example, and in many others, simple mean subtraction performs better and is thus the method used in our applications to correct for DC offsets.

In [46], the wrapped phase samples are retrieved as follows:

$$\varphi_i = \arctangent \left[\frac{1}{N} \sum_{n=0}^{N-1} S_B(iT + nT_s) e^{-j2\pi f_b n T_s} \right]. \quad (3.27)$$

This approach is mathematically equivalent to our approach of applying the Range FFT, retrieving the range bin signal and then performing arctangent demodulation. Equation 3.27 performs all three steps at once. Our approach assumes the target remains within the same range bin throughout. In Equation 3.27, this is equivalent to fixing the beat frequency, f_b , for all i . However, by not explicitly extracting the range bin signal, this approach makes it impossible to perform the magnitude analysis above and to analyse and correct for DC offsets. Correcting for DC offsets can especially improve the resulting phase signal as can be seen in Figure 3.5. For this reason, we prefer our step-by-step implementation over the all-at-once implementation.

For two of the three applications presented here, we are specifically interested in the frequency content of the radar phase signal. This can be obtained by applying the FFT on the radar phase signal. If the FFT is applied on the entire radar phase signal, then a single frequency estimate is obtained. With Figure 3.6 below, we demonstrate how a time-frequency analysis of the phase signals yields a time indexed frequency estimation. In the figure is the frequency of vibration of a speaker placed 70 cm from the radar. The frequency of the input signal into the speaker's amplifier is then ramped from 5 Hz at the start of the experiment, up to 60 Hz, and then down to 5 Hz again. In general, we used the small motions exhibited by the diaphragm of an 8 Ohm speaker as a testing model to verify the various parts of the theory presented here before conducting experiments on participants.

Frequency of the vibration of a speaker during a frequency sweep




Figure 3.6: It is possible to estimate a time indexed sequence of frequency estimates using our approach. This is achieved by applying a Short Time Fourier Transform (STFT) on the phase signals. Here, a time-indexed frequency estimation, from the radar phase signal, of the frequency of vibration of a speaker is illustrated. Specifically, this is experimental data from phase STFT results.

Chapter 4

Heart Rate Estimation with Radar

4.1 Introduction

In the previous chapter we demonstrated how applying the Range FFT to appropriate radar data reveals how far a target of interest is from the radar. We also demonstrated how this information can be used to extract the radar phase signal from the data. Additionally, we argued why this phase signal encodes information about the small motion of the target. In this chapter, we present the first of three applications of this theory. Specifically, we monitor the human heart rate non-invasively and without contact with more than 96% accuracy. Radar-based heart rate estimation is favourable because it is low-cost, non-contact and can be performed with clothing on.

4.2 Experimental Design

Throughout this entire project, six participants were recruited to provide data for the evaluation of our small motion monitoring techniques. For each of the three applications (heart rate estimation, respiration rate estimation and muscle activity monitoring), data to evaluate the technique is collected from some subset of these six participants. We refer to these participants as participant 1 to 6. For heart rate estimation, data was collected from participants 3, 4, 5 and 6. Three of these participants are male and one is female. All participants are aged between 22 and 25 years.

Four experiments were performed on each participant. We therefore have 16 experiments worth of data. Each experiment was about 52 seconds long and involved the participant sitting upright and still facing the radar in front of their chest. All participants were fully clothed and dressed as they arrived for these experiments i.e. no effort was made to influence the clothing of the participants. Figure 5.1 presents the radar rig used to collect data. The tripod on which the radar is mounted was adjusted such that the rod on top of which the radar and cameras sit (see Figure 5.1) was 87 cm from the ground. Each participant sat 70 cm from the tripod on a chair. The seat of the chair was adjusted to a height of 60 cm above the ground.

An effort was made to ensure that the radar points on the area of the chest that moves as a result of heart beat as much as possible. This area can easily be determined by placing a hand just below the left-hand side breast. This is important because the chest movement due to heart beat is about one to two orders of magnitude less than the movement due to respiration and can therefore be easily masked by this much larger movement.

For the duration of each experiment, each participant wears Maxim Integrated's MAXREFDES103# smart watch around their right-hand side wrist. The MAXREFDES103# uses a photoplethysmography (PPG) sensor to monitor heart rate. It achieves this by making use of a set of LEDs and a photodetector to measure blood volume changes [54]. We use the PPG heart

rate estimates as the reference or ground truth against which our radar heart rate estimates are compared. The radar data is sent to a desktop via ethernet and the PPG data is sent to the same desktop via USB. Both data are timestamped with the desktop’s time.

In all three of our applications, time-division multiplexing (TDM) was used to achieve a multiple-input-multiple-output (MIMO) system. The TI IWR6843ISK FMCW radar has three transmitters (Tx) and four receivers (Rx), allowing a total of 12 virtual receivers in TDM-MIMO [44]. With TDM-MIMO, only one of the three Tx is transmitting a chirp at any given time enabling the Rx to separate the signals from different Tx [48]. MIMO is useful for such things as angle-of-arrival estimation. In this work, we take advantage of the redundancy provided by the 12 virtual Rx to select the Rx channel with the best radar phase signal and process that signal. We noticed that the quality of the signals in the Rx channels often differed.

4.3 Results

Figure 4.1a below depicts the radar phase signal from one of the 16 experiments. This phase signal encodes the small motion of the chest due to heart beat and respiration as discussed in Chapter 3. The obvious periodic variation in the signal reflects respiration which is dealt with in the next chapter. The variation due to heart beat, however, is subtle and imperceptible to the eye.


Crudely, the respiration rate could be estimated simply by counting the number of cycles in the radar phase signal in Figure 4.1a. This is not possible for heart rate estimation. Instead, we rely on the Fourier transform and knowledge of the range of frequencies at which heart rate is to be expected. Figure 4.1b depicts the frequency spectrum of the radar phase signal in Figure 4.1a. The spectrum has a peak at 0 bpm due to stationary objects in the range bin of interest e.g. the bulk of the participant’s body. There is also a peak at about 22 bpm due to respiration and another one at about 44 bpm, which is a harmonic of respiration [20] [16] [55] [56].

As mentioned previously, the normal resting heart rate for an adult lies between 60 bpm and 100 bpm. In this work, we estimate the heart rate as the frequency corresponding to the peak in the range 60 – 100 bpm. In Figure 4.1b, this frequency, and thus the estimated heart rate, is 75.1 bpm. Normal healthy individuals can have heart rates less than 60 bpm; this is typical of athletes and was observed in participant 5. Moreover, other individuals can have heart rates above 100 bpm, for example during and after smoking [57] and such high heart rates were observed in participant 6. We therefore actually looked for the heart rate in the range 48 – 138 bpm. In [47] and [48], the same goal is achieved by passing the radar phase signal through a bandpass filter with cut-off frequencies at 48 bpm and 120 bpm.

The frequency spectrum in Figure 4.1b is obtained by taking the Fourier transform of the entire 52 seconds long radar phase signal in Figure 4.1a. We may, however, also be interested in how the heart rate varies within the 52 seconds long window. In order to represent the estimated heart rate in this way, we instead apply the STFT on the radar phase signal. Figure 4.2 depicts the resulting radar heart rate signal as a function of time as well as the PPG ground truth heart rate signal. The STFT was implemented by taking the FFT of 1.1 s successive segments of the radar phase signal. To improve the time resolution, 95% of the samples of each segment were overlapped with the next segment. A segment of 1.1 s duration is long enough for a variation in frequency of at least 0.91 Hz or 54.6 bpm. This duration should therefore be long enough time for the minimum frequency of interest of 60 bpm.


As a measure of the performance of our radar heart rate estimation approach, we calculated the error between the radar’s heart rate estimates and the PPG’s heart rate estimates. Specifically, we calculated both the root mean square error (RMSE) and the mean absolute error or average absolute error (AAE). Figure 4.3a presents the average RMSE and the AAE for each of our four participants

Radar phase signal encoding chest movements



(a) Radar phase signal encoding chest movements

Frequency spectrum of radar phase signal encoding heart rate



(b) Frequency spectrum of the radar phase signal encoding heart rate

Figure 4.1: The radar phase signal encodes the small motion of the chest due to heart beat and respiration. The obvious periodic variation in the signal reflects respiration while the variation due to heart beat, however, is subtle and imperceptible to the eye. We rely on the Fourier transform and knowledge of the range of frequencies at which heart rate is to be expected to estimate heart rate. For the data presented here, our heart rate estimate is 75.1 bpm.

Heart rate signals from PPG and radar




Figure 4.2: Heart rate signals from PPG (ground truth) and radar

averaged over the four experiments for each participant. The AAEs in Figure 4.3a were calculated as follows:

$$AAE = \frac{1}{E} \sum_{i=1}^E AAE_i = \frac{1}{E \cdot M} \sum_{i=1}^E \sum_{m=1}^M |HR_{rm} - HR_{PPGm}| \quad (4.1)$$

where E is the total number of experiments for each participant, M is the total number of samples in the radar heart rate signal, HR_{rm} is the m th sample of the radar heart rate signal and HR_{PPGm} is the m th sample of the PPG heart rate signal.

The average RMSEs in Figure 4.3a were calculated as follows:

$$RMSE = \frac{1}{E} \sum_{i=1}^E RMSE_i = \frac{1}{E} \sum_{i=1}^E \sqrt{\frac{1}{M} \sum_{m=1}^M (HR_{rm} - HR_{PPGm})^2}. \quad (4.2)$$

Additionally, we also calculated the accuracy of our heart rate estimates. Figure 4.3b presents both the mean and the median percentage accuracy of the radar heart rate estimates for each of the four participants averaged over the four experiments for each participant. The mean or average heart rate accuracy (AHRA) was calculated as follows:

$$AHRA = \frac{1}{E \cdot M} \sum_{i=1}^E \sum_{m=1}^M \left[100 - \left[\frac{|HR_{rm} - HR_{PPGm}|}{HR_{PPGm}} \times 100 \right] \right]. \quad (4.3)$$

The average median heart rate accuracy (MHRA) was calculated as

$$MHRA = \frac{1}{E} \sum_{i=1}^E MHRA_i, \quad (4.4)$$

Table 4.1: Comparison of the heart rate estimation accuracy of our work to others.

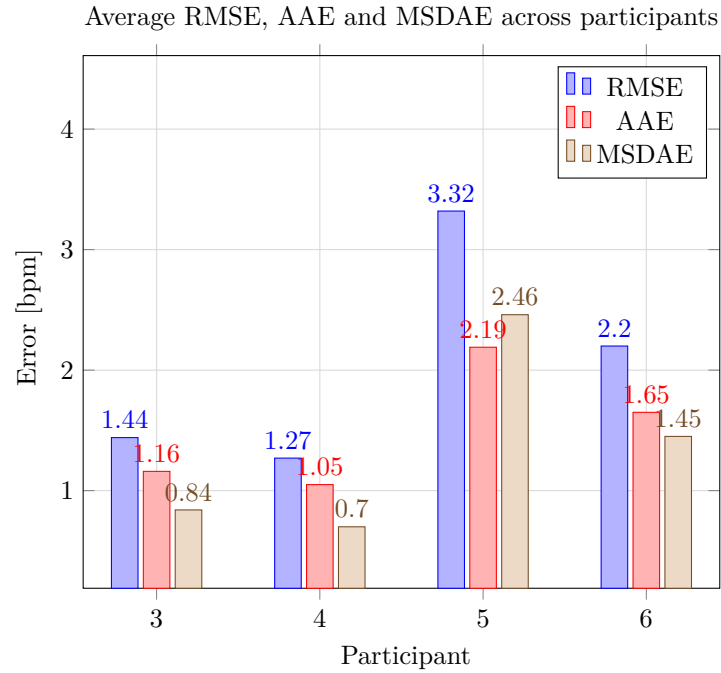
Reference	Distance (m)	f_{min} (GHz)	Accuracy (%)
Ours	0.7	60	98.5
[47]	1.7	77	80
[48]	1	77	99.2
[16]	1	5.46	98.5
[58]	2.48	9.6	55.2
[59]	1	80	87.2

where $MHRA_i$ is the median heart rate accuracy of experiment i .

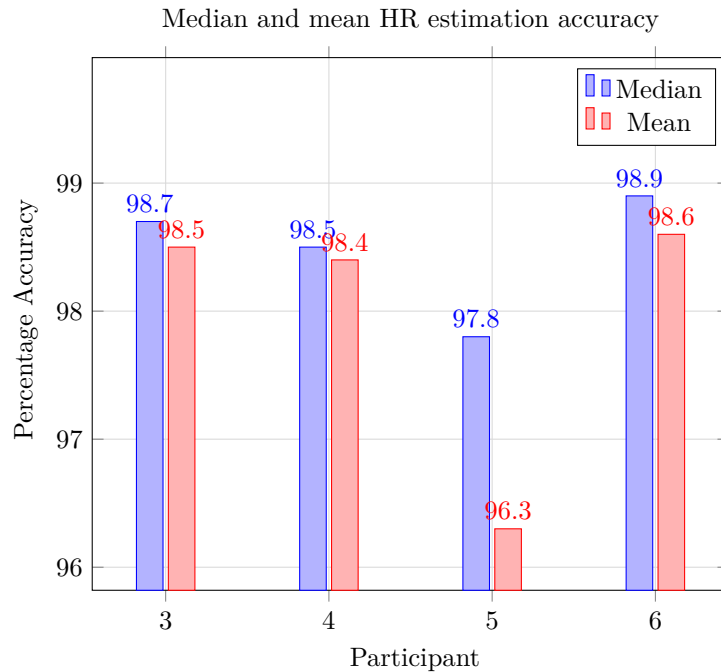
In all the participants except participant 5, our radar heart rate estimation approach achieves RMSE and AAE much less than 3 bpm. Our approach also achieves mean and median percentage accuracy greater than 98%, except in participant 5. We suspect that this poorer performance seen in participant 5 may be due to inaccurate ground truth estimates rather than or in addition to inaccurate radar estimates. This is based on the fact that PPG is known to suffer degraded performance on darker skin colours [11], which participant 5 has. Indeed, the MAXREFDES103# also reports how confident it is about its heart rate estimates as a percentage and it reported confidence values less than 85% much more frequently for participant 5.

The overall average RMSE, averaged over all four participants, is 2.06 bpm and the overall AAE is 1.51 bpm. Our overall median heart rate estimation accuracy is 98.5% and the overall mean accuracy is 98%. To make these results more informative, we also calculated the mean standard deviation of the absolute errors (MSDAE) from the average absolute errors for each participant. The results are also presented in Figure 4.3 below. The standard deviation value of 0.84 bpm for participant 3 means that 68% of the time, the absolute error between the radar and PPG heart rate estimates was within 0.84 bpm of 1.16 bpm. In other words, 68% of the time the absolute error between the radar and PPG heart rate estimates was between 0.32 bpm and 2.02 bpm for participant 3. The results for the other participants can be interpreted similarly. Finally, the overall standard deviation of the absolute errors, averaged over all four participants, is 1.36 bpm.

In Table 4.1 we compare the heart rate estimation accuracy of our approach to similar radar heart rate estimation approaches in the literature. For our approach and that in [16], the reported accuracy is the median accuracy. We also include how far the participant is from the radar and the starting frequency of the FMCW radar used. Our approach performs comparably to the best approaches and better than half of the approaches. It is to be noted, however, that we had the participants closest to the radar and performance tends to degrade the further a participant is from the radar.



(a) Radar HR estimation RMSE, AAE and MSDAE



(b) Radar HR estimation median and mean accuracy

Figure 4.3: The absolute errors between radar and PPG heart rate signals were calculated as a measure of performance. For each participant, these errors were averaged over all samples and experiments to give the average absolute errors. Additionally, the root mean square error, a common metric in non-contact heart rate estimation, was also calculated and both metrics are presented in (a). Also depicted in (a) are the average standard deviations of the absolute errors from the average absolute errors. From the average absolute errors, the accuracy of the radar heart rate estimation was calculated. Presented in (b) are the median and mean of these accuracies averaged over all samples and experiments for each participant.

Chapter 5

Respiration Rate Estimation with Radar and Cameras

5.1 Introduction

In Chapter 3 we discussed how the phase signal can be used to estimate small motions. Chapter 4 expanded upon this to demonstrate how we can monitor heart rate. Phase analysis allows for the estimation of motions so small that it would not be possible to estimate them through range alone. Specifically, these are motions smaller than a range bin. Take for example the movement of the chest due to breathing, with maximum displacement of about 4 – 12 mm [19] and our mmWave FMCW radar with range resolution of 37.5 mm at best (see Chapter 3). With this motion and range resolution, the chest would remain within the same range bin or the radar would register the same range estimate throughout.

In this chapter, as was demonstrated for heart rate monitoring, an analysis of the phase signal allows a very accurate estimate of the respiration rate. Additionally, we establish that, in conjunction with radar, breathing rate can be estimated with a stereo camera pair with very similar accuracy. Using triangulation, we use the two cameras to mimic the radar phase response and produce a camera phase signal. Admittedly, the \$732 cost of the GoPro HERO 7 Black stereo camera pair setup is comparable to the \$640 combined cost of the IWR6843ISK mmWave radar and DCA1000EVM radar data capture board. Moreover, the radar can penetrate through clothing [6], can operate in the dark and raises no privacy concerns.

However, cameras are ubiquitous, very easy to use and are already used for many purposes. Secondly, it is not necessary to use GoPros. The low frequency range of respiration means that almost any frame rate would be sufficient and a fine image resolution is unnecessary since the chest is broad and any point on it will do. Finally, unlike other camera-based respiration rate estimation approaches (for example motion capture [17]), our approach makes use of only two cameras compared to four, and one marker (which is very useful but not necessary) compared to 32 to 89.

5.2 Camera Based Phase Estimation

Our respiration rate monitoring experimental setup consists of the Texas Instruments (TI) IWR6843ISK mmWave FMCW radar and the DCA1000EVM data capture board. Both devices are enclosed within a 3D-printed casing which is then mounted on a 35 cm long rectangular 3D-printed rod. On either side of the radar is a GoPro HERO 7 Black camera (see Figure 5.1 below). We refer to these as camera 1 and 2. From the front, camera 1 is to the right of the radar while camera 2 is to the left. Both cameras are also mounted on the same rod. Finally, the rod is fixed to an adjustable height tripod.




Figure 5.1: Our respiration rate monitoring experimental setup consists of the Texas Instruments (TI) IWR6843ISK mmWave FMCW radar and the DCA1000EVM data capture board. Both devices are enclosed within a 3D-printed casing which is then mounted on a 35 cm long rectangular 3D-printed rod. On either side of the radar is a GoPro HERO 7 Black camera.

The radar’s phase signal, s , is obtained as per the discussion in Chapter 3, with the generic target replaced by a person’s chest. As a result, the phase signal encodes the small motion of the chest due to respiration. Here we discuss how this phase signal is computed using the image sequences (or videos) from cameras 1 and 2.

Using the theory of the relativistic Doppler effect [19], we can study how the frequency, f , of electromagnetic (EM) waves emitted by a source moving at velocity v relative to an observer changes. The perceived frequency change is often called the Doppler frequency shift, $f_D(t)$. This frequency shift can be quantified as $\frac{fv}{c}$. In a radar setup such as in Figure 5.1, the EM waves travel from the radar to the target and back. Therefore, the Doppler frequency shift observed at the radar due to the relative motion between the radar and the target is [19]

$$f_D(t) = \frac{2fv_r(t)}{c} = \frac{2}{\lambda}v_r(t), \quad (5.1)$$

where $v_r(t)$ is the radial component of the target’s velocity (for a stationary radar).

The frequency of the signal received by the radar will be shifted by a value equal to $f_D(t)$ [19].

We also know that the instantaneous Doppler frequency shift due to some motion should be related to the phase signal due to that same motion through

$$f_D(t) = \frac{1}{2\pi} \frac{d\varphi(t)}{dt}. \quad (5.2)$$

Equating equations 5.1 and 5.2 and integrating to solve for the phase gives

$$\varphi(t) = \frac{4\pi}{\lambda} \int_0^t v_r(\tau) d\tau. \quad (5.3)$$

If we can approximate $v_r(t)$ from the video data, then we can estimate the phase signal via equation 5.3. We now consider how to estimate the radial component of the target's velocity, $v_r(t)$, from video. Figure 3.3 has been reproduced as Figure 5.2 below with the inclusion of cameras 1 and 2 in green (reproduced notation holds the same meaning). In black is the radar's coordinate system with origin at O while camera 1's coordinate system ($x_c - y_c - z_c$) is in yellow with origin at O_c . Suppose that the target is within the fields of view of both cameras throughout the entire time the radar is transmitting and the cameras recording. Also suppose that the target moved from point A to point B in a time equal to $T_c = \frac{1}{f_c}$, where f_c is the frame rate of both cameras. We refer to $\vec{V}_t = \dot{\vec{X}}_t$ as the instantaneous velocity of the target at point A.




Figure 5.2: Our respiration rate estimation setup consists of a stereo camera pair and a radar between the cameras. Suppose that the target is within the fields of view of both cameras and the radar. Using triangulation, we use the two cameras to mimic the radar and produce a camera phase signal.

Our task is to resolve \vec{V}_t along \vec{R}_0 . Let us approximate \vec{V}_t as $\vec{V}_t = \frac{\vec{X}_t}{T_c}$. Notice that we can quantify \vec{X}_t if we know the coordinates of both points A and B in the same coordinate system. We calibrate cameras 1 and 2 by using an A4 checkerboard pattern to obtain the intrinsic and extrinsic parameters. By identifying point A in an image from camera 1 and in the corresponding image from camera 2, we can triangulate to obtain the 3D coordinates of point A with respect to the yellow frame. The same procedure can be repeated for point B. We now have vectors \vec{O}_cA and \vec{O}_cB , respectively. Finally, we have $\vec{X}_t = \vec{O}_cB - \vec{O}_cA$ and $\vec{V}_t = \frac{\vec{O}_cB - \vec{O}_cA}{T_c}$.

To resolve this \vec{V}_t along \vec{R}_0 , we must now compute \vec{R}_0 . It is easy to see that $\vec{R}_0 = \vec{O}\vec{O}_c + \vec{O}_cA$, where $\vec{O}\vec{O}_c$ is the position of O_c relative to O and \vec{O}_cA is represented in the black frame. One of the extrinsic parameters we get from stereo camera calibration is the position of camera 2 in the yellow frame. Let this position be $(L, 0, 0)$. Then $\vec{O}\vec{O}_c = (-\frac{L}{5}, 0, 0)$. Note that this assumes that the

radar lies exactly on the black x-axis on a point between camera 1 and 2, one-fifth of the way from camera 1. Given our hardware (see Figure 5.1), this is a reasonable assumption. Finally, we must transform the O_cA we get from triangulation from the yellow frame to the black frame. By observing what happens to the three unit vectors \hat{i} , \hat{j} and \hat{k} of the yellow frame under the transformation, T , we are looking for, we find

$$T = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}. \quad (5.4)$$

Generally, at any time $t = lT_c$ ($l \in \{0, 1, 2, \dots, F - 1\}$, where F is the total number of video frames) with the target at point P , we have the velocity vector \vec{V}_t and the line of sight (LOS) $\vec{OP} = O\vec{O}_c + O_c\vec{P}$. We then find the radial component of the target's velocity by taking the dot product of \vec{V}_t with the unit vector $\frac{O\vec{P}}{|OP|}$.

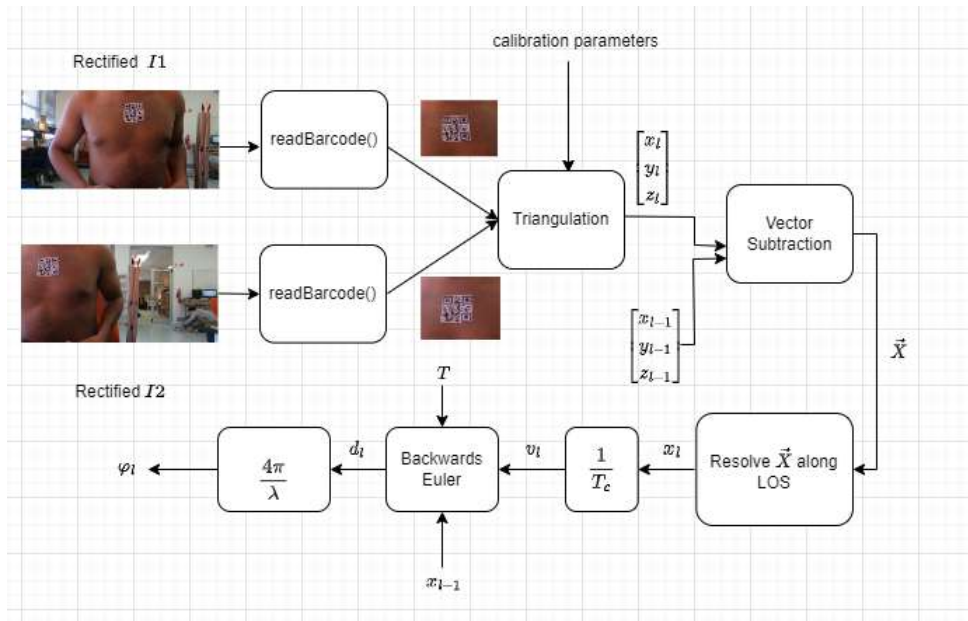


Figure 5.3: A rectified stereo pair of images is passed through `readBarcode()` to obtain the 2D pixel coordinates of a point of interest (POI). These coordinates are used together with calibration parameters to compute the 3D world coordinates of the POI. This way, instantaneous displacements can be computed for all times t , $0 \leq t \leq F \cdot T_c$.

Figure 5.3 summarises the camera-based phase estimation for respiration rate monitoring. The two images of the chest from cameras 1 and 2, I1 and I2 respectively, are first rectified (or undistorted). Specifically, the target is a point on the chest. We place a QR code on the chest and ensure that one of the three finder patterns is placed over this point. This allows us to make use of the `readBarcode()` method from MATLAB's Computer Vision Toolbox, which returns the 2D pixel coordinates of the finder patterns and thus the point of interest in the image. With these coordinates and the calibration parameters, we triangulate to obtain the 3D world coordinates of the point. Note that the QR code on the chest is unnecessary since one can track the point of interest manually by selecting the right pixel in each image. However, this is very tedious and time-consuming.

Figure 5.3 is just for illustrative purposes. In reality, all processing is done after the video data has been captured i.e. post processing. This allows us to smoothen the often noisy radial velocity

signal and to interpolate it so that the camera based phase has as many samples as the radar phase. This is why we use the chirp time T in our discrete integration (we use the backward Euler method) instead of the inter-image time T_c .

5.3 Experimental Design

We recruited four participants to provide data for the evaluation of our respiration rate monitoring technique. These are the same participants 1, 4, 5 and 6 from Chapter 4. All four participants were male between the ages of 22 and 25. If necessary for the QR code to stick, the skin on the chest was shaved with a brand new razor blade for each participant and then the QR code was pasted on the shaved skin patch with paper glue.

The tripod was adjusted such that the rod on top of which the radar and cameras sit was 87 cm from the ground. Each participant sat 70 cm from the tripod on a chair. The seat of the chair was adjusted to a height of 60 cm above the ground.

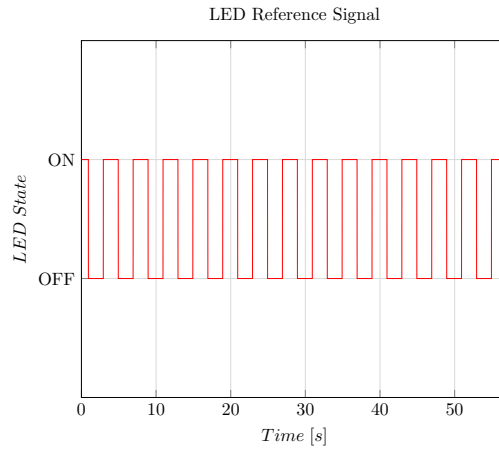
We take advantage of the fact that respiration can be controlled voluntarily. We asked each participant to breath at a known frequency. As a cue to help the participants breath at this frequency, we blink a big orange LED in front of them. The LED is driven by a signal generator producing a square wave at this frequency. When the LED is on, we ask the participants to breath in and when the LED is off, we ask them to breath out. A similar LED, driven by the same signal, is placed within the field of view of the cameras. By tracking this LED's brightness across video frames, we get a ground truth signal against which we compare our phase based radar and camera respiration signals. The LED's frequency was changed from experiment to experiment, taking on 200 mHz, 250 mHz and 300 mHz (that is 12 bpm, 15 bpm and 18 bpm respectively).

5.4 Results

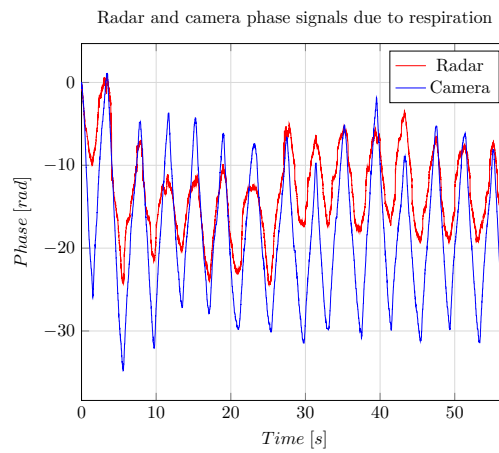
The results presented here are from a total of 16 experiments across the four participants. Four experiments were conducted for each participant. Each experiment is about 1 minute (57.1 s) long. Of these experiments, eight had the LED blinking at 200 mHz (12 bpm), three at 250 mHz (15 bpm) and five at 300 mHz (18 bpm). Figures 5.4 and 5.5 present the data from one of these experiments.

For the experiment presented in Figures 5.4 and 5.5, the LED was blinking at 250 mHz (15 bpm). Figure 5.4a presents the LED's ON or OFF state signal during the experiment. The corresponding radar and camera phase signals are presented in Figure 5.4b. It can be seen that whenever the LED's state transitions from ON to OFF, the phase signals change from decreasing to increasing. The participant changes from breathing in to breathing out during an ON to OFF transition. As a result of exhaling air from the lungs, their chest moves inward and away from the radar and therefore the displacement, and thus the phase, increases. The reverse is true for an OFF to ON transition.

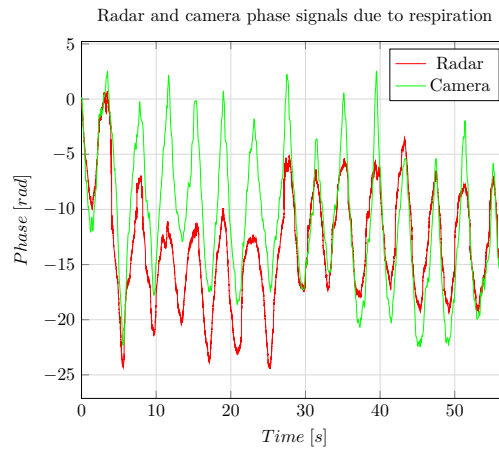
Recall from Section 5.2 that the camera phase signal is computed from the radial component of the chest's velocity. In addition to getting this component by resolving \vec{V}_t along the line of sight (LOS), \vec{OP} , at each time t (assuming the point of interest on the chest is at P), we also experimented with simply resolving \vec{V}_t along the radar frame's y-axis (see Figure 5.2) for all time. Of course, this y-axis only approximates the LOS and the approximation gets worse the further away point P is away from the y-axis. Since we are specifically interested in small motion, this approximation is reasonable. The advantage of this approximation is that the LOS does not have to be updated at each time t . For this reason, we say this LOS is fixed, as opposed to updated. Figures 5.4c and 5.5c present the results obtained using the fixed LOS to compute the camera phase signal.



(a) State of the reference LED during the experiment



(b) Phase signals encoding respiration (LOS updated)



(c) Phase signals encoding respiration (LOS fixed)

Figure 5.4: Radar and camera data recorded during one of the respiration rate experiments. (a) ON-OFF state of the reference LED during the experiment. (b) Recorded radar and camera phase signals encoding chest movement due to respiration (LOS updated). (c) Recorded radar and camera phase signals encoding chest movement due to respiration (LOS fixed).

Figure 5.4 depicts how the LED state, radar phase and camera phase change with time. Complementarily, Figure 5.5 illustrates the frequency content of these same signals. It can be seen that both the radar and camera phase frequency spectra have a dominant peak around 15 bpm as expected. To be exact, the radar phase’s peak is at 15.09 bpm, the camera phase’s peak at 15.12 bpm and the LED reference signal’s peak at 15.21 bpm. For the fixed LOS, the camera phase’s peak is at 15.09 bpm, which is the same as the frequency estimated by the radar phase. These frequency values corresponding to the peaks are the respiration rate estimates for the two sensors.

Over all 16 experiments, our non-contact respiration rate estimation system achieved average absolute errors (AAE) of 0.09 bpm and 0.10 bpm for the radar and camera, respectively. Using the fixed LOS approximation, the achieved AAE for the camera respiration rate estimation is 0.13 bpm. We calculate AAE for the radar respiration rate estimates as

$$AAE = \frac{1}{E} \sum_{i=1}^E |RR_{ri} - RR_{*i}|, \quad (5.5)$$

where E is the total number of experiments, RR_{ri} is the respiration rate estimated by the radar in experiment i and RR_{*i} is the ground truth respiration rate in experiment i .

AAE for the camera respiration rate estimates is calculated the same way with RR_{ri} replaced by RR_{ci} , which is the respiration rate estimated by the camera in experiment i .

Another metric often quoted in vital sign estimation literature is the respiration rate accuracy. Similarly to [60], we calculate the overall average respiration rate accuracy (ARRA) of our radar estimates as a percentage:

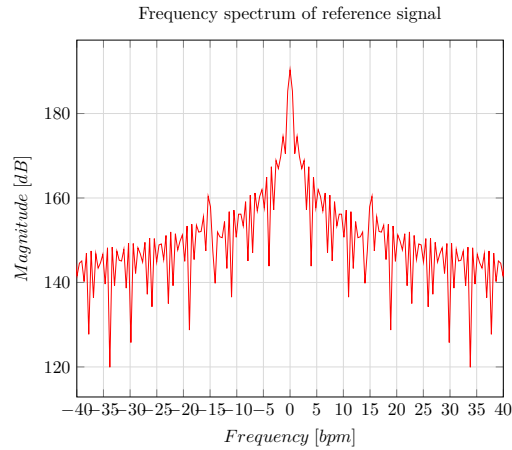
$$ARRA = \frac{1}{E} \sum_{i=1}^E \left[100 - \left[\frac{|RR_{ri} - RR_{*i}|}{RR_{*i}} \times 100 \right] \right]. \quad (5.6)$$

Just as it is for AAE, the ARRA for the camera respiration rate estimates is calculated by replacing RR_{ri} with RR_{ci} .

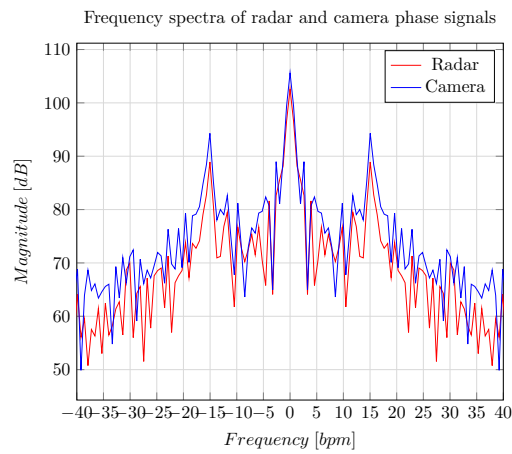
Our system achieved an ARRA of 99.35% for the radar estimates and an ARRA of 99.26% for the camera estimates. Using the fixed LOS approximation, the achieved ARRA for the camera respiration rate estimation is 99.06%. Based on the reported AAEs and ARRAs above, notice that the fixed LOS approximation achieves very slightly poorer performance as compared to the updated LOS phase computation.

Table 5.1 presents a comparison of the respiration rate estimation accuracy achieved in our work to other work available in the literature. Only the work in [48] achieves better performance than ours. Granted, our participants were closer to the radar than in other work though not by much. In this work, we estimated the respiration rate by taking the Fourier Transform of the entire 57.1 s long radar phase signal. This gives us good frequency resolution but poor time resolution since we get only one respiration rate estimate per 57.1 s, with no access to how the respiration rate varies within the 57.1 s.

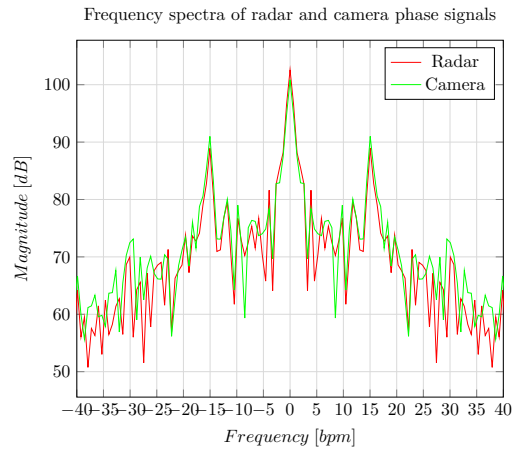
We also investigated how the AAE changes with or is affected by each of our four participants and how it changes with or affected by what the ground truth respiration rate is. The bar graphs in Figure 5.6 present both the radar and camera (LOS updated) AAE for each participant and for each ground truth respiration rate. Though admittedly, we experimented on a few participants and distinct ground truth respiration rates, no marked change is observed on AAE across both participants and ground truth respiration rates.



(a) Frequency spectrum of the LED reference signal



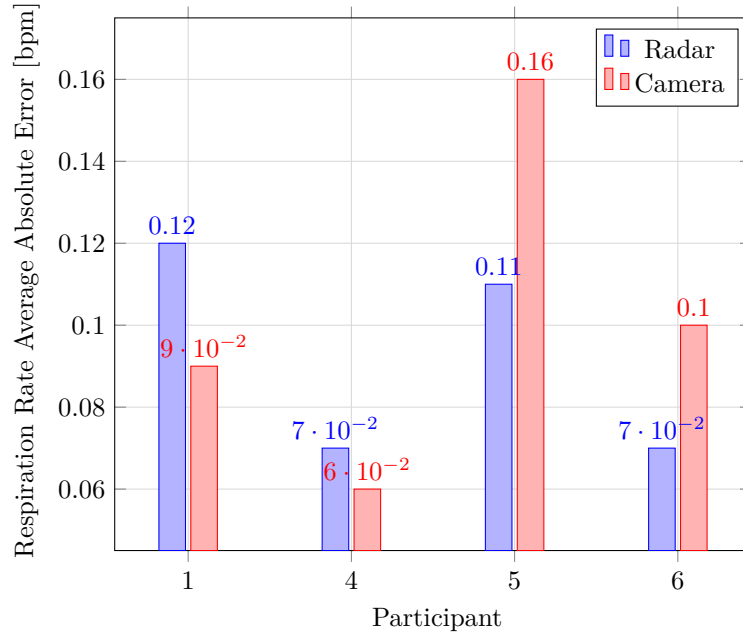
(b) Frequency spectra of the phase signals (LOS updated)



(c) Frequency spectra of the phase signals (LOS fixed)

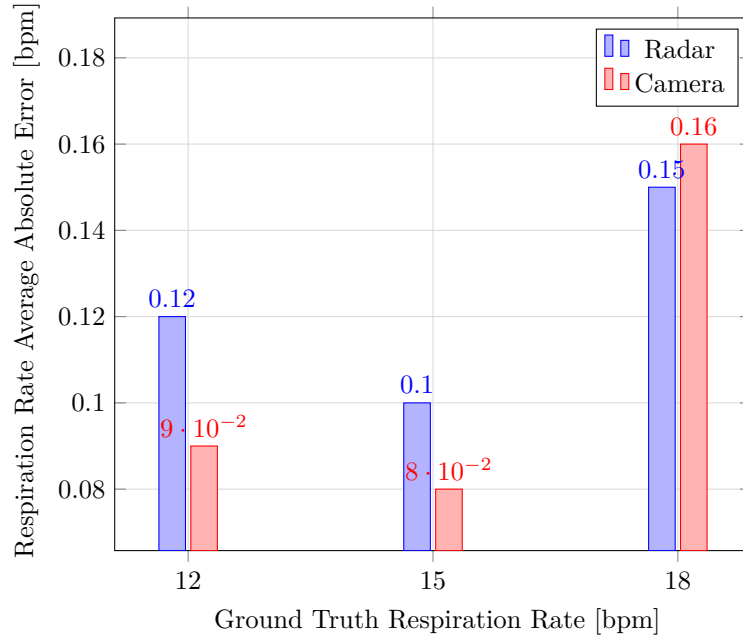
Figure 5.5: Frequency spectra of the radar and camera data recorded during one of the respiration rate experiments. (a) Frequency spectrum of the LED reference signal. (b) Frequency spectra of the radar and camera phase signals encoding chest movement due to respiration (LOS updated). (c) Frequency spectra of the radar and camera phase signals encoding chest movement due to respiration (LOS fixed).

Radar and camera respiration rate estimation average absolute errors across participants



(a) Respiration rate absolute errors across participants

Radar and camera respiration rate estimation average absolute errors across respiration rates



(b) Respiration rate estimation errors across respiration rates

Figure 5.6: Presented here are the average absolute errors of our respiration rate estimates averaged over experiments for each participant. We then averaged these errors across experiments and participants to give the average errors across the ground truth respiration rates.

Table 5.1: Comparison of the respiration rate estimation accuracy of our work to others.

Reference	Distance (m)	Accuracy (%)
Ours	0.7	99.4
[47]	1.7	94
[48]	1	100
[16]	1	99.3
[59]	1	91.1

Recall that the camera phase signal estimates the radar phase signal. Because of this, we also investigated how well the camera phase signal estimates the radar phase signal. We do this by calculating the Pearson’s correlation coefficient r , the average absolute error, the standard deviation of the average absolute error and the slope m and bias c of the best fit line between the two phase signals. The Pearson’s correlation coefficient is a measure of how two datasets are similar. It assumes a value between -1 and 1. Values 1 and -1 imply that the two datasets are perfectly positively correlated and perfectly negatively correlated, respectively. A value of 0 implies that the datasets are uncorrelated. We calculate this value as [47]

$$r = \frac{\sum_{i=1}^M (\varphi_{ri} - \overline{\varphi_r})(\varphi_{ci} - \overline{\varphi_c})}{\sqrt{\sum_{i=1}^M (\varphi_{ri} - \overline{\varphi_r})^2} \sqrt{\sum_{i=1}^M (\varphi_{ci} - \overline{\varphi_c})^2}}, \quad (5.7)$$

where φ_{ri} is the i th radar phase signal sample, φ_{ci} is the i th camera phase signal sample, $\overline{\varphi_r}$ is the average of the radar phase signal, $\overline{\varphi_c}$ is the average of the camera phase signal and M is the total number of received chirps.


For the same experiment whose results are presented in Figures 5.4 and 5.5, the Pearson’s correlation coefficient between the radar phase signal and the camera phase signal (LOS updated) is 0.83. Using the fixed LOS approximation, the correlation coefficient is 0.74. Over all the 16 experiments, the average Pearson’s correlation coefficient is 0.62 (LOS updated) and 0.69 for the fixed LOS approximation. Both of these values are considered moderate positive correlation, with the 0.69 for the fixed LOS approximation indicating borderline strong positive correlation [61].

In Figure 5.7 are scatter plots of the camera phase signal against the corresponding radar phase signal. These are the same phase signals in Figure 5.4. Also presented in these plots are the lines along which the dependent variable is equal to the independent variable i.e “ $y = x$ ” line or where the camera phase signal is equal to the radar phase signal in our case. Finally, the lines that best fit the phase data are also depicted. If the camera phase signal was a perfect estimate of the radar phase signal (that is, $r = 1$), then the phase samples, the line that best fits them and the “ $y = x$ ” line would all coincide. Note that departure from the ideal cannot all be accounted for in the camera phase signal. The radar phase signal itself is prone to errors as discussed in Section 3.6.

For small values of the radar phase (that is, close to 0), we see that the best fit line is fairly close to the “ $y = x$ ” line. For larger values of the radar phase, the lines diverge. The small motions that interest us correspond to small values of the radar phase and thus the camera phase signal should be a reasonable estimate. The best fit line in Figure 5.7a (LOS updated) has slope $m = 1.29$ and bias $c = -1.08$ rad while that in Figure 5.7b (LOS fixed) has slope $m = 0.88$ and bias $c = 0.69$ rad. In the ideal case, we would have $m = 1$ and $c = 0$ rad.

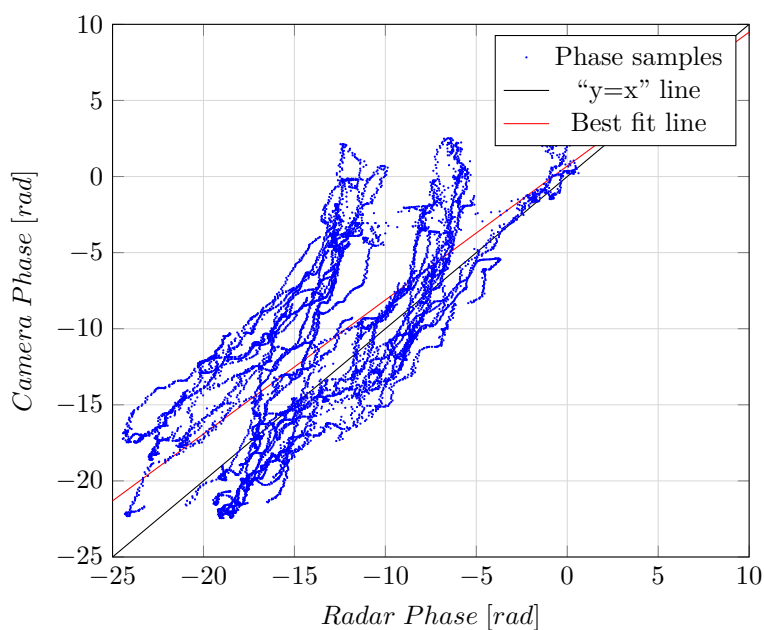
Over the 16 experiments, the average slope of the best fit line is $m = 0.80$ (LOS updated) and $m = 0.76$ (LOS fixed). The slope is a measure of how sensitive the camera phase is to changes in the underlying small motion, relative to the radar phase. A value of the slope less than 1 indicates that the camera phase is less sensitive to the underlying small motion compared to the radar phase.

Scatter plot of camera phase versus radar phase (LOS updated)



(a) Scatter plot of camera phase versus radar phase (LOS updated)

Scatter plot of camera phase versus radar phase (LOS fixed)



(b) Scatter plot of camera phase versus radar phase (LOS fixed)

Figure 5.7: Camera phase signal versus radar phase signal scatter plots illustrating the relationship between the two signals. If the camera phase signal was a perfect estimate of the radar phase signal, then the phase samples, the line that best fits them and the " $y = x$ " line would all coincide. (a) Scatter plot illustrating the relationship of camera phase to radar phase (LOS updated). (b) Scatter plot of camera phase data against radar phase data (LOS fixed).

As one would expect, our average slope values indicate that the camera phase is less sensitive than the radar phase. Additionally, our average slope values indicate that the LOS updated approach is more sensitive than the fixed LOS approach. This is also evident in Figure 5.4.

The average bias of the best fit line is $c = -5.71$ rad (LOS updated), which translates to a bias of 2.27 mm in the underlying small motion. Using the fixed LOS approximation, the average bias is $c = -4.96$ rad, which translates to a bias of 1.98 mm in the underlying small motion. In the frequency estimation, the bias appears as a component at DC (at 0 Hz). The lower bound of normal respiration rates (≈ 0.2 Hz) is large enough for this DC component, and thus the bias, to not be a problem for our camera based respiration rate estimation.

In this study, we used the brightness signal from an LED blinking at a known frequency as a cue for our participants to follow as they breathe normally. Though we did not have access to a conventional respiration rate sensor (as we did for heart rate, see Chapter 4) that could give us respiration rate ground truth across time to compare our estimates to, it is possible to estimate a time indexed sequence of respiration rates using our approach. This is achieved by performing a time-frequency analysis, for example by applying a Short Time Fourier Transform (STFT) on the phase signals instead of a Fourier transform on the entire phase signal all at once.

Chapter 6

Muscle Activity Monitoring with Radar and Cameras

6.1 Introduction

EMG, though it is expensive, is vital for rehabilitation from neurological and orthopaedic disorders [23]. It monitors muscle activity by measuring the electrical signals from the brain that excite muscles [21]. Other methods of monitoring muscle activity, such as laser doppler myography (LDMi) and Sonomyography (SMG), exist. LDMi is a non-contact approach, exploiting the fact that muscles vibrate when active [25]. LDMi measures the minute displacement of the skin over the vibrating muscle [42] while SMG measures the shape change or deformation of the muscle [37]. EMG, however, is still the most widely used muscle activity monitoring approach [41]. Unfortunately, the “gold standard” EMG system, Noraxon, is expensive at more than \$29,000. Additionally, electrodes can be uncomfortable and invasive (needle electrodes) to patients and are often prone to electrical noise (surface electrodes).

In this chapter we demonstrate that, using the same hardware and algorithms from the respiration rate estimation chapter, we can monitor muscle activity using radar phase and camera phase signals. For respiration rate monitoring, the phase signals encode the small motion exhibited by the surface of the skin on the chest due to respiration. Here, the phase signals encode the small motion exhibited by the surface of the skin over the muscle of interest due to muscle activity.

In [33], Jun Shi et al. used ultrasound images to track the motion of the surface of the biceps brachii muscle during an isometric contraction. From their ultrasound images, it is clear that this muscle’s surface is just underneath the skin surface. It is also clear that the skin surface, whose motion we monitor here, moves with the muscle’s surface. They draw the conclusion that their SMG signal could potentially be used to complement EMG as a muscle assessment tool.

Similarly, in [37], optical flow is used to track the thickness of the brachioradialis muscle. It is demonstrated that this muscle deformation correlates with elbow output force at a level comparable to EMG during different isometric contractions. Our phase based muscle activity monitoring approach measures the small motion at the surface of the skin over the muscle of interest. This motion is due to the muscle’s motion and/or deformation, measured in [33] and [37], respectively. Because of this, our approach measures the same signal measured by sonomyography (SMG). We call it radiomyography (RMG).

6.2 Experimental Design

To evaluate the performance of our radar and camera based muscle activity monitoring system, we recruited three participants from which the required data was collected. These are participants 1, 5 and 6 from Chapters 4 and 5 above. From each participant, during each experiment, three sets of time-synchronised data were collected. These are radar, video and EMG data. The radar and EMG sensors were externally triggered to start recording by a common signal generated by a signal generator. When a button on the signal generator is pressed to send the trigger signal, a light turns on in the scene of the cameras. This light is then used to synchronise the cameras with the other two sensors.

Four experiments, each about 1 minute (57.1 s) long, were conducted on each participant. During an experiment, a participant was asked to isometrically contract and relax their vastus lateralis muscle (the strongest and largest quadriceps muscle that runs along the entire length of the lateral side of the thigh) randomly. As with the respiration rate estimation, a QR code was pasted on the skin over the vastus lateralis muscle with paper glue. If necessary for the QR code to stick, the skin was shaved with a brand new razor blade for each participant. A pair of EMG electrodes was placed over the vastus lateralis and one EMG ground electrode was placed over the patellar tendon (see Figure 6.1).

Each participant sits on top of a table that is placed 70 cm in front of the tripod while wearing shorts to expose the QR code to the cameras, as presented in Figure 6.1. The tripod's height is adjusted such that the radar's receivers and transmitters are at the same height as the participant's vastus lateralis muscle. An effort is made to ensure the vastus lateralis muscle is on the radar's boresight.




Figure 6.1: Each participant sits on top of a table that is placed 70 cm in front of the tripod while wearing shorts to expose the QR code to the cameras. Also in the figure is the signal generator used to provide a common trigger signal to the radar and EMG sensor. The bright blue light on one of the signal generator's buttons is used to synchronise the cameras with the other two sensors.

6.3 Results

The results presented here are from a total of 12 experiments across the three participants. Four experiments were conducted for each participant. Each experiment is about 1 minute (57.1 s) long.

Figure 6.2 below illustrates the EMG, radar phase and camera phase signals recorded from one such experiment, specifically from participant 1. We observe that each time the participant contracts the vastus lateralis, the radar and camera phase signals increase. This implies that contraction deforms the vastus lateralis muscle away from the radar-camera rig. Conversely, when the muscle is relaxed, the radar and camera phases decrease. This suggests that relaxation deforms the muscle towards the radar-camera rig.

The most noticeable feature of Figure 6.2 is the high correlation between the EMG signal and the deformation signals (as measured by the phase signals). Several other studies [33] [34] [37] have demonstrated that muscle deformation (as measured through ultrasound images) correlates with muscle activity (as measured through EMG) and/or muscle force.

Another relationship that was observed between normalised EMG and normalised radar (and camera) phase is that during contraction, the deformation is an exponential function of the EMG. The signals are normalised by dividing each sample by the largest sample in the signal. This exponential relationship was also observed in [33] and [62] where deformation was measured through SMG. Figure 6.3b presents a scatter plot of normalised radar phase against normalised EMG for the first contraction in Figure 6.2. The exponential relationship between the two variables can be seen in this plot.

Also depicted in Figure 6.3b is an exponential curve that best fits the data. Following [33] and [62], the best fit line was modelled as

$$Y = A(1 - e^{-BX}), \quad (6.1)$$

where Y is the normalised muscle deformation, X is the normalised EMG, A is the asymptotic value of Y and B is the exponent coefficient determining the curvature.


To solve for the model coefficients A and B , the problem was posed as a non-linear least-squares curve fitting problem and solved using the Levenberg–Marquardt algorithm. For the data in Figure 6.3b, the values of A and B were found to be 0.91 and 3.68, respectively. The equation of the depicted best fit line is thus $Y = 0.91(1 - e^{-3.68X})$.

The model coefficients were calculated for all 12 experiments. Figure 6.4 below presents the 12 values of each coefficient across participants and experiments. The average values of A were found to be 1, 0.91 and 1.06 for participants 1, 5 and 6 respectively. For B , the average values were found to be 3.45, 12.83 and 3.60 for participants 1, 5 and 6 respectively. Based on [33] and [62], these are typical values. In [33], the average value of B for the biceps brachii muscle over seven participants is 6.19. In [62] it is 7.69 for the biceps brachii muscle over five participants whereas we have 6.63 as the average value of B over all three participants.

In Figure 6.4 we present the values of A and B for each participant separately. This is because the value of B may be fundamentally specific to each participant. Indeed, in [33] it was established through ANOVA that there was a significant difference in B across the seven participants. Differences in B across participants mean that the percentage change in muscle deformation per percentage change in the driving EMG varies across participants. The relatively high average B value of 12.83 for participant 5 (see Figure 6.4b) suggests that participant 5 has the highest percentage change in muscle deformation per percentage change in the driving EMG. Of course, this then means that the muscle deformation for this participant saturates at lower percentage values of the EMG.


We do not calculate the Pearson’s correlation coefficient between radar (or camera) phase and EMG because the relationship is not linear. In fact, we have established that it is exponential (see Figure 6.3b). Now that has been established, we measure the goodness of fit of our exponential model to the observed data. We do this by computing the coefficient of determination R^2 . The

EMG signal from left vastus lateralis muscle



(a) EMG signal from left vastus lateralis muscle

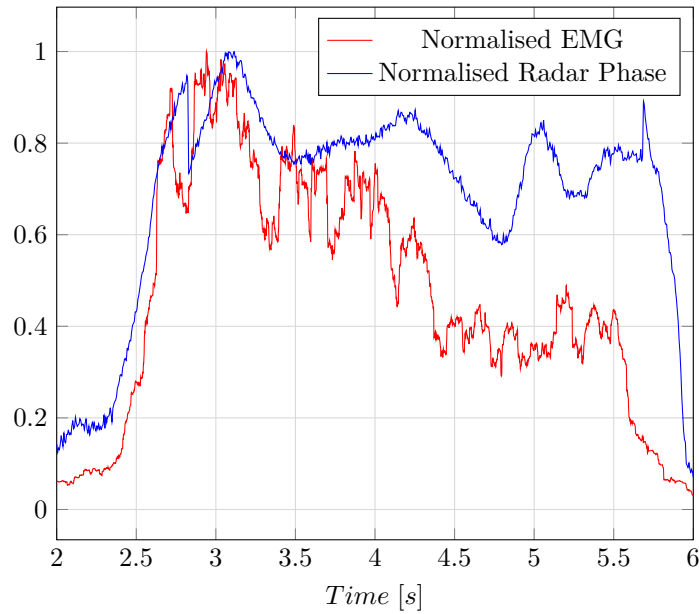
Radar phase and camera phase signals



(b) Radar phase and camera phase signals


Figure 6.2: During each experiment, a participant was asked to isometrically contract and relax their vastus lateralis muscle randomly. Three sets of time-synchronised data were then collected. These are radar, video and EMG data as presented here.

Normalised radar phase and EMG over a contraction-relaxation cycle




(a) Normalised EMG and normalised radar phase

Normalised radar phase versus normalised EMG from the vastus lateralis




(b) Normalised radar phase versus normalised EMG and best fit line

Figure 6.3: Illustrated in (a) is the normalised EMG and corresponding normalised radar phase from a single contraction-relaxation cycle in Figure 6.2b. The relationship observed between the two normalised signals is characteristic of muscles and, represented differently, it results in the hysteresis curve in Figure 6.6. A scatter plot of the normalised radar phase against normalised EMG during the contraction phase of (a) is presented in (b). Also included in (b) is the exponential line that best fits the data. The exponential relationship observed here is typical.



(a) Value of coefficient A across participants and experiments



(b) Value of coefficient B across participants and experiments

Figure 6.4: The relationship of normalised radar phase to normalised EMG during the contraction of a muscle has been observed to be exponential, as depicted in Figure 6.3b. Equation 6.1 is often used to model the relationship and we do the same here. In (a) is the value of coefficient A and in (b) is the value of coefficient B for the data collected from all the experiments viewed per participant. MATLAB's `lsqnonlin()` method from the Optimisation Toolbox is used to fit Equation 6.1 to the data.

coefficient of determination is a measure of the proportion of the variation in the dependent variable (the muscle deformation) that is predictable from the independent variable (the EMG) using our exponential model. The value of R^2 is calculated as follows:

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - f_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}, \quad (6.2)$$

where n is the total number of samples in the data used, y_i is the i th sample of the observed normalised radar phase, \bar{y} is the mean value of the observed normalised radar phase and f_i is the i th sample of the normalised radar phase predicted by the model.

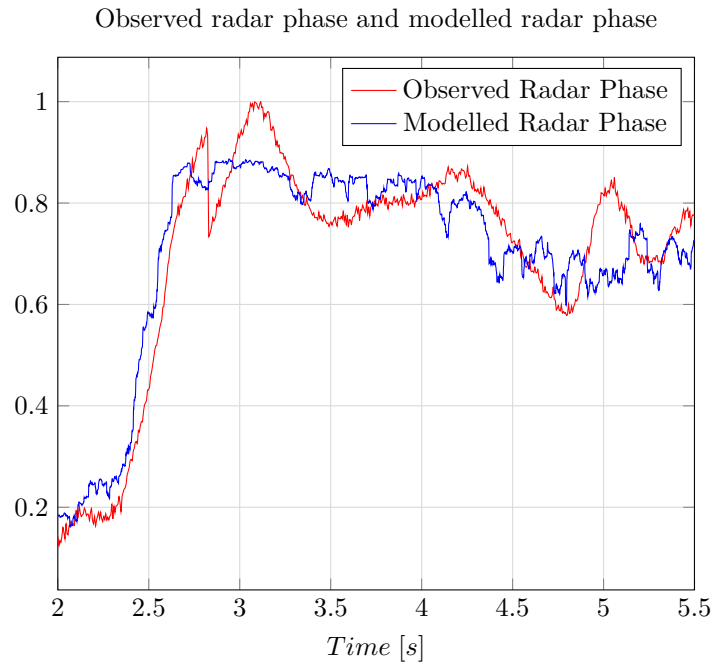
Equation 6.2 gives a value between 0 and 1. The goodness of fit for the illustrative experiment in Figure 6.3 is 0.88. This means that 88% of the muscle deformation is predicted or accounted for by the EMG. Figure 6.5b presents the average value of R^2 for each of the three participants, averaged over the four experiments. We found the average value of R^2 across all participants to be 72% whereas in [33] they found the average value of R^2 to be 88%.

Another interesting phenomenon or relationship between EMG and muscle deformation can be observed in Figure 6.3a. The plots in this figure can be considered in two stages, the ON or contraction stage from about 2 s to about 3 s and the OFF or relaxation stage from about 5.5 s to about 6 s. Notice that a given percentage of muscle deformation (normalised radar phase) corresponds to a lower percentage EMG during the OFF stage than during the ON stage. This phenomenon is easily observable when the normalised EMG is plotted against the normalised radar phase for only the ON and OFF stages. Figure 6.6 below depicts such a plot, drawn from the data in Figure 6.3a. The resultant curve is called a hysteresis curve and was also observed in [33] and [63]. In [33], the hysteresis is between muscle deformation and EMG and is exactly as we observe here. In [63] however, it is observed between muscle deformation and force and it is stated that such hysteresis curves are well known in biomechanics. EMG increases during the ON stage and decreases during the OFF stage, and therefore the hysteresis in Figure 6.6 (and [33]) has a clockwise direction while that in [63] has an anti-clockwise direction.


The final observation common between our study and the studies in [33] and [63] is that the rate of muscle deformation is typically higher during the ON stage compared to the OFF stage. As in [33] however, this effect was not always observed. It was almost exclusively observed only in participant 1. In [33], this limitation was attributed to the low frame rate (8 Hz) of their ultrasound imaging system and the inability of the participants to reduce torque smoothly. In our study, the muscle deformation was sampled at the slow time frequency of 178.5 Hz, which is more than 20 times that in [33]. However, no effort was made to ensure that the muscle was relaxed smoothly.

Additionally, it was noticed that the data from participant 1 contained significantly less random body motion (RBM) as compared to participants 5 and 6. RBM is any motion of the body that is not the motion of interest. This RBM may also be the reason why participants 5 and 6 had lower average values of the coefficient of determination than participant 1 (see Figure 6.5b). This is plausible since that would mean some of the variation in the muscle deformation that is unaccounted for by EMG is accounted for by the RBM.

Finally, another limitation of our approach is that the recovered muscle deformation is prone to drift. This means that after contraction, the radar phase does not return back to 0 rad even though the muscle has been relaxed (see Figure 6.2b). This happens because after every contraction, the muscle does not return to exactly the same position relative to the radar-camera rig as it was before the contraction. Though an effort was made by each participant to avoid translation of the leg, the leg was not strapped in place. The SMG approach is also prone to this drift problem as demonstrated in [34], [35] and [37]. Ways of correcting this drift for ultrasound recovered muscle deformation exist. It was hypothesised in [37] that drift was the cause of diminished performance and it was demonstrated in [34] that correcting this drift can improve some performance measures.



(a) Observed and modelled radar phase



(b) Coefficient of determination between EMG and radar phase

Figure 6.5: Using Equation 6.1 and the model parameters in Figure 6.4, we computed the normalised radar phase predicted by the model given some normalised EMG from our data. In (a) is the resulting normalised radar phase and the normalised radar phase actually observed from the experiment in Figure 6.2b. With these computed radar phases, we computed the goodness of fit of our exponential models for all experiments. In (b) is the goodness of fit across participants.

Normalised EMG versus normalised radar phase during ON and OFF stages

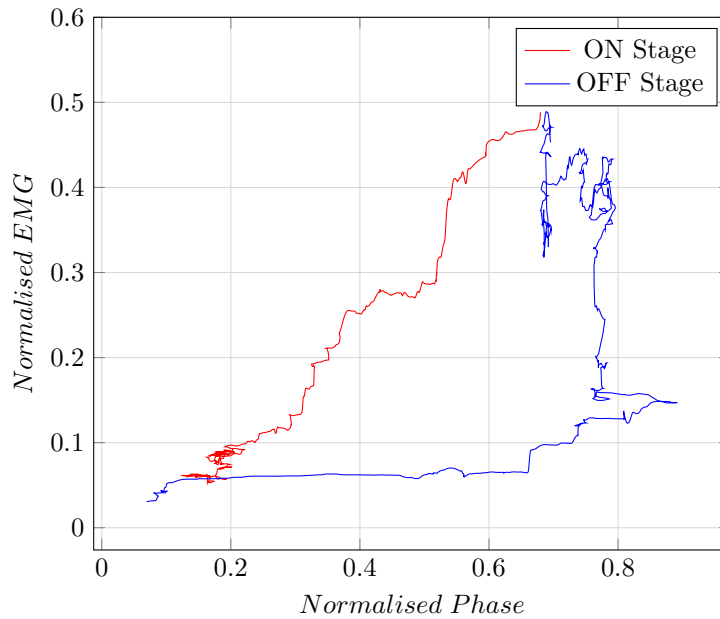


Figure 6.6: Presented here is the normalised EMG data from the illustrative experiment and illustrative contraction-relaxation cycle plotted against the corresponding normalised radar phase. Only the ON and OFF stages of the contraction-relaxation cycle are presented; the part of the cycle when the contraction is held is not depicted. The hysteresis curve observed here is characteristic and well known in biomechanics.

Similarly, it was observed in our study that contraction stages severely affected by drift have very poor coefficients of determination.

Based on the foregoing results and discussions, we argue that the recovered radar phase signal is the same muscle deformation signal recovered through SMG. Several studies [33] [34] [35] [37] [63] have demonstrated that this signal correlates with muscle force. More importantly, experiments performed in [37] have demonstrated that participants can be instructed to use this signal to track a trajectory. This deformation-based trajectory tracking can be performed with accuracy similar to that achieved with EMG-based trajectory tracking and participants prefer using this deformation signal. The muscle deformation signal is thus a new avenue for device control, e.g. the control of prosthetics.

However, unlike SMG, our approach is not only non-invasive but also non-contact. SMG is technically challenging to implement without contact because of the difference between the acoustic impedance of air and that of human tissue. At the interface between media with different acoustic impedances, some of the acoustic energy is transmitted into the other medium while some is reflected [39]. To reduce the amount of reflected energy at the air-skin interface, the acoustic impedance of the skin must match that of the media the ultrasound transducer is in (i.e. the two media must be acoustically coupled). To achieve this, an ultrasound gel is often applied on the skin and the transducer placed in it [33], making it impossible for the transducer to be ‘non-contact’ in any real sense.


All the muscle deformation signals presented above and the accompanying discussions have been from and about radar phase signals. However, we perform respiration rate and muscle activity monitoring using both radar and cameras independently. Figure 6.7a below presents the mean Pearson’s correlation coefficient between the radar phase signal and the camera

phase signal across participants, averaged over the four experiments for each participant. One can see that the radar phase signal and the camera phase signal are highly correlated for all participants.

Because of the similarity between the radar phase signal and the camera phase signal, the relationships observed between EMG and radar phase are also observed between EMG and the camera phase signal. Figure 6.7b demonstrates that the same exponential relationship observed between normalised EMG and normalised radar phase is also observed between normalised EMG and normalised camera phase. In addition, Figure 6.7c demonstrates that the same hysteresis effect observed between normalised EMG and normalised radar phase is also observed between normalised EMG and normalised camera phase.


Recall from Chapter 5 on respiration rate estimation, we found that updating the LOS resulted in slightly better performance compared to keeping the LOS fixed for the camera-based respiration rate estimation. However, for the muscle activity monitoring experiments in this chapter, updating the LOS lead to much worse camera-based results compared to keeping the LOS fixed. One major difference between the experimental setups for respiration rate and muscle activity is that the QR code in the muscle activity experiments is much lower in the video frames than it is in the respiration rate experiments (compare Figures 5.3 and 6.1). Recall again that if a target moves back and forth on the y-axis in Figure 5.2, then its LOS to the radar remains fixed. We suspect then that by virtue of the QR code (the “target”) being closer to the y-axis in the muscle activity experiments than it is in the respiration rate experiments, assuming that the LOS is fixed is a better approximation than whatever errors are incurred in the computations that update the LOS (see Section 5.2).

Mean radar-camera phases pearson's correlation coefficient across participants



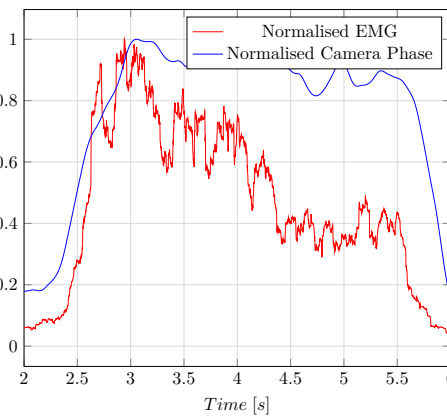
(a) Mean radar-camera phases pearson's correlation coefficient

Normalised Camera Phase versus Normalised EMG from the Vastus Lateralis



(b) Normalised camera phase versus normalised EMG

Normalised Camera Phase and Normalised EMG



(c) Normalised camera phase and normalised EMG

Figure 6.7: (a) Presents the mean Pearson's correlation coefficient between the radar phase signal and camera phase signal across participants, averaged over the four experiments for each participant. (b) Depicts that the same exponential relationship observed between normalised EMG and normalised radar phase is also observed between normalised EMG and normalised camera phase. (c) Illustrating the relationship between normalised EMG and normalised camera phase, which is also the same hysteresis effect observed between normalised EMG and normalised radar phase.

Chapter 7

Limitations, Recommendations and Conclusions

7.1 Introduction

In this chapter we discuss important limitations of our non-contact small-motion monitoring approach. After highlighting each limitation, we recommend a potential way of addressing it in future work. Finally, conclusions from the work are drawn.

7.2 Limitations and Recommendations

7.2.1 Random Body Motion (RBM) Cancellation

A challenge that faces any approach that monitors an object's small motion is the presence of larger motions that corrupt the signal of interest. For this reason, our approach requires the participant to remain still during the data recording process. The radar signal phase in Figure 4.1a can be seen to slowly drift upwards during the experiment. This is due to the participant's torso leaning away from the radar during the experiment. If this RBM was sufficiently faster and of sufficient magnitude, then the respiration or heart rate of interest would be masked.

We have already seen how the larger amplitude movements of the chest due to respiration mask the component of the signal from the smaller amplitude movements due to heart beat. If the frequency of the corrupting signal is sufficiently larger or smaller than the frequency of the signal of interest, filtering can solve this challenge. This is the approach that most work in the literature take to extract heart rate from the radar signal phase, in the presence of only respiration. We also take a similar approach. Filtering is, however, not reasonable when the frequency ranges of the signals are unknown or are very close.

As discussed in Chapter 6, RBM causes drifts in the resulting radar (and camera) signal phase. These drifts are not much of a problem for frequency analysis and thus respiration and heart rate estimation. For muscle activity monitoring however, where we are interested in the degree of activity, drifts are an issue (see Figure 6.2a).

A potential way to achieve random body motion cancellation (RBMC) is to augment the small motion monitoring setup with a separate sensor that just senses the RBM. This can be achieved, for example, by using a camera as the RBM sensor while the radar measures both RBM and small motion. For muscle activity monitoring, specifically for the vastus lateralis muscle, the camera can measure just RBM if the QR Code is placed on the knee. Subtracting the

camera phase signal from the radar phase signal will result in a signal encoding just the small motion.

With Figure 7.1 and 7.2 below, we demonstrate that this can potentially work. Figure 7.1 shows the Doppler frequency shift due to a speaker vibrating at 60 Hz while being moved back and forth to introduce RBM over its vibratory small motion. The RBM starts at 5 s and it completely masks the 60 Hz vibration. Figure 7.2 shows the Doppler shift after RBMC. Unlike in Figure 7.1, the 60 Hz component can be seen for all time after RBMC. Figure 7.3 depicts the experimental setup used to demonstrate this RBM cancellation. The platform on which the speaker is mounted is moved by hand to introduce the RBM. The two cameras are used to measure just the RBM, using the aforementioned QR Code approach, which is then subtracted from the radar phase signal.




Figure 7.1: Micro-Doppler signature of speaker vibrating at 60 Hz with random body motion (RBM). The RBM is introduced by moving the speaker back and forth randomly.

7.2.2 Radar Camera Calibration

In Chapter 5 we discussed how the camera data can be related to the radar data for our rig. However, this relation relied on geometric approximations of where the radar and cameras were relative to each other. In future work, a data-driven approach, similar to [64] and [65] can be implemented for a more accurate radar-camera calibration.

7.2.3 Signal Quality Assessment

In Chapter 3 we discussed that the radar phase signal is prone to artefacts when the range bin signal has a very small magnitude. Artefacts in the radar phase signal degrade performance in our applications of interest. Equation 3.23 tells us that the range bin signal will have a small magnitude if the received chirp echo has a small magnitude. In [6], the radar is placed on a servo platform which pans the radar until it faces the human in an orientation that gives optimal radar signal quality.

In future work, we recommend similarly moving the radar-camera rig to improve performance of our system. This would lead to huge benefits as we observed that the quality of the received chirp echoes was very sensitive to the relative radar-participant orientation especially for heart rate estimation, as so did [6]. However, instead of using the Peak To Average (PTA) metric as a measure of radar phase signal quality as in [6], we propose using the sigmoid of the magnitude of the range bin

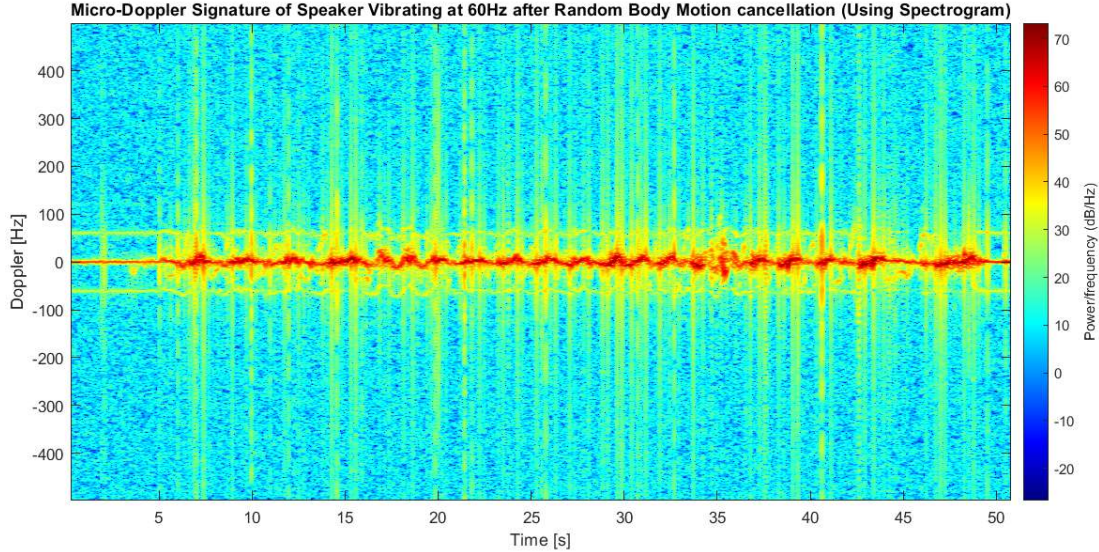


Figure 7.2: Micro-Doppler signature of the speaker after RBM cancellation. The cancellation is achieved by subtracting a camera phase signal from the radar phase signal. The camera phase signal encodes only RBM while the radar phase signal encodes both RBM and small motion.

signal as a measure of radar phase signal quality. We observed this metric to be a reliable measure of the quality of the radar signal.

7.2.4 Participant Tracking

Our small motion monitoring approach assumes that the target remains within one range bin. This is reasonable for the motions we are interested in here. However, we may also want to monitor these small motions for participants moving around freely.

For future work, we recommend extending our approach to enable the tracking and monitoring of small motions for mobile participants. To achieve this, we propose not only estimating the target's range only for the first received chirp echo but estimating the range on every received chirp and then looking in the appropriate range bin for the phase sample. Alternatively, if the varying range of the target is known, say from another range sensor, then the radar phase signal can be computed using Equation 3.27 with the beat frequency, f_b , updated using the range information.

7.2.5 Robust Heart Rate Estimation

In Chapter 4, on heart rate estimation with radar, we mentioned that the frequency associated with the breathing rate had harmonics, as also observed in [16], [20], [55] and [56]. These harmonics can be misconstrued as the frequency corresponding to the heart rate.

In the future, this challenge could be overcome by devising a model that would enforce some temporal continuity in the heart rate estimates. For example, a Hidden Markov Model (HMM) could be used to track continuity or capture temporal dependencies in the sequential data.

7.3 Conclusions

The current state of the art for monitoring heart rate is electrocardiography (ECG), while that for monitoring muscle activity is electromyography (EMG). Both are contact-based approaches that are uncomfortable, expensive, time-consuming, and require a good understanding of anatomy,




Figure 7.3: A speaker playing a 60 Hz tone is placed on top of a wooden platform. Pasted on the platform is a QR Code. The platform is moved back and forth by hand to corrupt the vibratory motion perceived by the radar with RBM. The two cameras measure just the RBM by making use of the QR Code.

for example, for electrode placement. Our non-contact approach addresses all these challenges. Furthermore, radar and cameras have complementary strengths. Radar can penetrate clothing, is insensitive to lighting conditions, and raises no privacy concerns. On the other hand, cameras are ubiquitous, very easy to use, and already used for many purposes.

In this dissertation, we presented the theory of a radar- and camera-based small motion monitoring approach with clear mathematical models and justifications. By complementarily using radar and cameras in the way we did, we believe other interesting problems can be explored. We demonstrated how the single theory we presented can be used to estimate heart rate and respiration rate and monitor muscle activity without contact. In all our applications, we achieve performance comparable to the current state of the art. To the best of the author's knowledge, our approach is the first to measure the characteristic dimensional changes of muscles in vivo and without contact. Finally, we highlighted the important limitations of our work, together with potential ways to address them in the future.

Bibliography

- [1] G. Balakrishnan, F. Durand, and J. Guttag, “Detecting pulse from head motions in video,” in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 3430–3437.
- [2] R. E. Klabunde, *Cardiovascular Physiology Concepts*, 2nd ed. Philadelphia, PA: Lippincott Williams and Wilkins, Sep. 2011.
- [3] A. J. Pappano and W. Gil Wier, “The arterial system,” in *Cardiovascular Physiology*. Elsevier, 2013, p. 2.
- [4] G. J. Betts, J. Wise, K. A. Young, P. Desaix, E. W. Johnson, J. E. Johnson, O. Korol, D. Kruse, B. Poe, and M. D. Womble, *Anatomy and physiology*. XanEdu Publishing Inc, 2017.
- [5] “Pqrst.” [Online]. Available: <https://www.medicalexamprep.co.uk/the-basics-of-ecg-interpretation-part-2-rate-rhythm-and-axis/pqrst/>
- [6] P. Zhao, C. X. Lu, B. Wang, C. Chen, L. Xie, M. Wang, N. Trigoni, and A. Markham, “Heart rate sensing with a robot mounted mmWave radar,” in *International Conference on Robotics and Automation (ICRA)*, 2020.
- [7] J. Lin, J. Kiernicki, M. Kiernicki, and P. Wollschlaeger, “Microwave apexcardiography,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 27, no. 6, pp. 618–620, 1979.
- [8] H.-Y. Wu, M. Rubinstein, E. Shih, J. Guttag, F. Durand, and W. Freeman, “Eulerian video magnification for revealing subtle changes in the world,” *ACM Trans. Graph.*, vol. 31, no. 4, pp. 1–8, Aug. 2012.
- [9] W. Wang, A. C. den Brinker, S. Stuijk, and G. de Haan, “Algorithmic principles of remote PPG,” *IEEE Trans. Biomed. Eng.*, vol. 64, no. 7, pp. 1479–1491, Jul. 2017.
- [10] X. Wang, C. Yang, and S. Mao, “PhaseBeat: Exploiting CSI phase data for vital sign monitoring with commodity WiFi devices,” in *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, Jun. 2017.
- [11] E. M. Nowara, D. McDuff, and A. Veeraraghavan, “A meta-analysis of the impact of skin type and gender on non-contact photoplethysmography measurements,” in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020, pp. 1148–1155.
- [12] M. M. Cloutier, *Respiratory physiology*, 2nd ed., ser. Mosby’s Physiology Monograph. Philadelphia, PA: Elsevier - Health Sciences Division, Oct. 2018.
- [13] H. Abdelnasser, K. A. Harras, and M. Youssef, “UbiBreathe,” in *Proceedings of the 16th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. ACM, Jun. 2015. [Online]. Available: <https://doi.org/10.1145/2746285.2755969>
- [14] A. Davies and C. Moores, *The respiratory system*, 2nd ed., ser. Systems of the Body. Churchill Livingstone, Feb. 2014.

- [15] I. Ajayi, “Limbic modulation of the autonomic nervous system,” Ph.D. dissertation, The University of Queensland, 2017.
- [16] F. Adib, H. Mao, Z. Kabelac, D. Katabi, and R. C. Miller, “Smart homes that monitor breathing and heart rate,” in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, Apr. 2015.
- [17] C. Massaroni, A. Nicolò, M. Sacchetti, and E. Schena, “Contactless methods for measuring respiratory rate: A review,” *IEEE Sensors Journal*, vol. 21, no. 11, pp. 12 821–12 839, 2021.
- [18] A. Singh, S. U. Rehman, S. Yongchareon, and P. H. J. Chong, “Modelling of chest wall motion for cardiorespiratory activity for radar-based NCVS systems,” *Sensors (Basel)*, vol. 20, no. 18, p. 5094, Sep. 2020.
- [19] V. C. Chen, *The micro-Doppler effect in radar the micro-Doppler effect in radar*, 2nd ed. Norwood, MA: Artech House, Feb. 2019.
- [20] D. R. Morgan and M. G. Zierdt, “Novel signal processing techniques for Doppler radar cardiopulmonary sensing,” *Signal Processing*, vol. 89, no. 1, pp. 45–66, Jan. 2009.
- [21] R. Merletti and D. Farina, Eds., *Surface Electromyography*, ser. IEEE Press Series on Biomedical Engineering. Nashville, TN: John Wiley & Sons, Apr. 2016.
- [22] C. Joseph, A. Delcarme, I. Vlok, K. Wahman, J. Phillips, and L. Nilsson Wikmar, “Incidence and aetiology of traumatic spinal cord injury in cape town, south africa: a prospective, population-based study,” *Spinal Cord*, vol. 53, no. 9, pp. 692–696, Sep. 2015.
- [23] R. H. Chowdhury, M. B. I. Reaz, M. A. B. M. Ali, A. A. A. Bakar, K. Chellappan, and T. G. Chang, “Surface electromyography signal processing and classification techniques,” *Sensors (Basel)*, vol. 13, no. 9, pp. 12 431–12 466, Sep. 2013.
- [24] A. Manca, A. Cereatti, L. Bar-On, A. Botter, U. Della Croce, M. Knafitz, N. A. Maffioletti, D. Mazzoli, A. Merlo, S. Roatta, A. Turolla, and F. Deriu, “A survey on the use and barriers of surface electromyography in neurorehabilitation,” *Front. Neurol.*, vol. 11, p. 573616, Oct. 2020.
- [25] C. Orizio, “Muscle sound: bases for the introduction of a mechanomyographic signal in muscle studies,” *Crit. Rev. Biomed. Eng.*, vol. 21, no. 3, pp. 201–243, 1993.
- [26] M. Tarata, A. Spaepen, and R. Puers, “The accelerometer MMG measurement approach, in monitoring the muscular fatigue,” 2001.
- [27] M. T. Tarata, “Mechanomyography versus electromyography, in monitoring the muscular fatigue,” *Biomed. Eng. Online*, vol. 2, p. 3, Feb. 2003.
- [28] R. B. Woodward, M. J. Stokes, S. J. Shefelbine, and R. Vaidyanathan, “Segmenting mechanomyography measures of muscle activity phases using inertial data,” *Sci. Rep.*, vol. 9, no. 1, p. 5569, Apr. 2019.
- [29] R. B. Woodward, S. J. Shefelbine, and R. Vaidyanathan, “Pervasive monitoring of motion and muscle activation: Inertial and mechanomyography fusion,” *IEEE/ASME Transactions on Mechatronics*, vol. 22, pp. 2022–2033, 2017.
- [30] N. Campbell, T. Egan, and C. Deegan, “The application of digital accelerometers for wired and non-wired mechanomyography,” *2017 28th Irish Signals and Systems Conference (ISSC)*, pp. 1–6, 2017.
- [31] J. W. Rohrbaugh, E. J. Sirevaag, and E. J. Richter, “Laser Doppler vibrometry measurement of the mechanical myogram,” *Rev. Sci. Instrum.*, vol. 84, no. 12, p. 121706, Dec. 2013.

- [32] M. O. Ibitoye, N. A. Hamzaid, J. M. Zuniga, N. Hasnan, and A. K. A. Wahab, "Mechanomyographic parameter extraction methods: an appraisal for clinical applications," *Sensors (Basel)*, vol. 14, no. 12, pp. 22 940–22 970, Dec. 2014.
- [33] J. Shi, Y.-P. Zheng, Q.-H. Huang, and X. Chen, "Continuous monitoring of sonomyography, electromyography and torque generated by normal upper arm muscles during isometric contraction: sonomyography assessment for arm muscles," *IEEE Trans. Biomed. Eng.*, vol. 55, no. 3, pp. 1191–1198, Mar. 2008.
- [34] L. A. Hallock, A. Velu, A. Schwartz, and R. Bajcsy, "Muscle deformation correlates with output force during isometric contraction," in *2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*, 2020, pp. 1188–1195.
- [35] A. T. Kamatham, M. Alzamani, A. Dockum, S. Sikdar, and B. Mukherjee, "A simple, drift compensated method for estimation of isometric force using sonomyography," in *Sensing Technology*, ser. Lecture notes in electrical engineering. Cham: Springer International Publishing, 2022, pp. 355–366.
- [36] L. Brausch, H. Hewener, and P. Lukowicz, "Classifying muscle states with one-dimensional radio-frequency signals from single element ultrasound transducers," *Sensors*, vol. 22, no. 7, 2022. [Online]. Available: <https://www.mdpi.com/1424-8220/22/7/2789>
- [37] L. A. Hallock, B. Sud, C. Mitchell, E. Hu, F. Ahamed, A. Velu, A. Schwartz, and R. Bajcsy, "Toward real-time muscle force inference and device control via optical-flow-tracked muscle deformation," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 29, pp. 2625–2634, Dec. 2021.
- [38] T. A. McMahon, *Muscles, reflexes, and locomotion*. Princeton, NJ: Princeton University Press, Apr. 1984.
- [39] P. Regtien and E. Dertien, "9 - acoustic sensors," in *Sensors for Mechatronics (Second Edition)*, 2nd ed., P. Regtien and E. Dertien, Eds. Elsevier, 2018, pp. 267–303. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/B9780128138106000094>
- [40] H. Azhari, "Special imaging techniques," in *Basics of Biomedical Ultrasound for Engineers*. Hoboken, NJ, USA: John Wiley & Sons, Inc., Apr. 2010, p. 313.
- [41] S. Casaccia, L. Scalise, L. Casacanditella, E. P. Tomasini, and J. W. Rohrbaugh, "Non-contact assessment of muscle contraction: Laser Doppler Myography," in *2015 IEEE International Symposium on Medical Measurements and Applications, MeMeA 2015, Torino, Italy, May 7-9, 2015*. IEEE, 2015, pp. 610–615. [Online]. Available: <https://doi.org/10.1109/MeMeA.2015.7145276>
- [42] C. Orizio, M. Gobbo, A. Veicsteinas, R. Baratta, B. Zhou, and M. Solomonow, "Transients of the force and surface mechanomyogram during cat gastrocnemius tetanic stimulation," *European Journal of Applied Physiology*, vol. 88, no. 6, pp. 601–606, Feb. 2003.
- [43] V. Dham, "Programming chirp parameters in TI radar devices," accessed: 31-08-2022. [Online]. Available: https://www.ti.com/lit/an/swra553a/swra553a.pdf?ts=1661930658263&ref_url=https%253A%252F%252Fwww.ti.com%252Fproduct%252FIWR6843AOP
- [44] S. Rao, "MIMO radar (rev. a) - Texas Instruments," Jul 2018. [Online]. Available: <https://www.ti.com/lit/an/swra554a/swra554a.pdf>
- [45] "60ghz mmwave sensor evms," May 2022. [Online]. Available: <https://www.ti.com/lit/ug/swru546e/swru546e.pdf?ts=1694578369717>
- [46] Y. Xiong, Z. Peng, G. Xing, W. Zhang, and G. Meng, "Accurate and robust displacement measurement for FMCW radar vibration monitoring," *IEEE Sensors Journal*, vol. 18, no. 3, pp. 1131–1139, 2018.

- [47] M. Alizadeh, G. Shaker, J. C. M. de Almeida, P. P. Morita, and S. Safavi-Naeini, "Remote monitoring of human vital signs using mmWave FMCW radar," *IEEE Access*, vol. 7, pp. 54 958–54 968, 2019.
- [48] Z. Xu, C. Shi, T. Zhang, S. Li, Y. Yuan, C.-T. M. Wu, Y. Chen, and A. Petropulu, "Simultaneous monitoring of multiple people's vital sign leveraging a single phased-MIMO radar," *IEEE Journal of Electromagnetics, RF and Microwaves in Medicine and Biology*, vol. 6, no. 3, pp. 311–320, 2022.
- [49] A. Davis, "Visual vibration analysis," Ph.D. dissertation, Massachusetts Institute of Technology, August 2016.
- [50] M. Özcan, "Differentiation method for phase recovery," in *Practical Holography XXIX: Materials and Applications*, vol. 9386. SPIE, 2015, pp. 41–48.
- [51] J. Wang, X. Wang, L. Chen, J. Huangfu, C. Li, and L. Ran, "Noncontact distance and amplitude-independent vibration measurement based on an extended DACM algorithm," *IEEE Transactions on Instrumentation and Measurement*, vol. 63, no. 1, pp. 145–153, 2014.
- [52] B.-K. Park, S. Yamada, and V. Lubecke, "Measurement method for imbalance factors in direct-conversion quadrature radar systems," *IEEE Microwave and Wireless Components Letters*, vol. 17, no. 5, pp. 403–405, 2007.
- [53] A. Singh, X. Gao, E. Yavari, M. Zakrzewski, X. H. Cao, V. M. Lubecke, and O. Boric-Lubecke, "Data-based quadrature imbalance compensation for a CW Doppler radar system," *IEEE Transactions on Microwave Theory and Techniques*, vol. 61, no. 4, pp. 1718–1724, 2013.
- [54] P. Kyriacou and J. Allen, *Photoplethysmography: Technology, Signal Analysis and Applications*. Elsevier Science, 2021. [Online]. Available: <https://books.google.co.uk/books?id=d8wnEAAAQBAJ>
- [55] M. Mercuri, I. R. Lorato, Y.-H. Liu, F. Wieringa, C. Van Hoof, and T. Torfs, "Vital-sign monitoring and spatial tracking of multiple people using a contactless radar-based sensor," *Nat. Electron.*, vol. 2, no. 6, pp. 252–262, Jun. 2019.
- [56] C. Li and J. Lin, "Random body movement cancellation in Doppler radar vital sign detection," *IEEE Transactions on Microwave Theory and Techniques*, vol. 56, no. 12, pp. 3143–3152, 2008.
- [57] S. Ramakrishnan, K. Bhatt, A. K. Dubey, A. Roy, S. Singh, N. Naik, S. Seth, and B. Bhargava, "Acute electrocardiographic changes during smoking: an observational study," *BMJ Open*, vol. 3, no. 4, p. e002486, 2013. [Online]. Available: <https://doi.org/10.1136/bmjopen-2012-002486>
- [58] L. Anitori, A. de Jong, and F. Nennie, "FMCW radar for life-sign detection," in *2009 IEEE Radar Conference*, 2009, pp. 1–6.
- [59] S. Wang, A. Pohl, T. Jaeschke, M. Czaplik, M. Köny, S. Leonhardt, and N. Pohl, "A novel ultra-wideband 80 Ghz FMCW radar system for contactless monitoring of vital signs," in *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2015, pp. 4978–4981.
- [60] J.-M. Lee, H. Song, and H.-C. Shin, "Respiration rate extraction of moving subject using velocity change in FMCW radar," *Appl. Sci. (Basel)*, vol. 12, no. 9, p. 4114, Apr. 2022.
- [61] B. Ratner, "The correlation coefficient: Its values range between +1/−1, or do they?" *J. Target. Meas. Anal. Mark.*, vol. 17, no. 2, pp. 139–142, Jun. 2009.
- [62] P. W. Hodges, L. H. M. Pengel, R. D. Herbert, and S. C. Gandevia, "Measurement of muscle contraction with ultrasound imaging," *Muscle Nerve*, vol. 27, no. 6, pp. 682–692, Jun. 2003.

- [63] C. Orizio, M. Gobbo, A. Veicsteinas, R. V. Baratta, B. H. Zhou, and M. Solomonow, “Transients of the force and surface mechanomyogram during cat gastrocnemius tetanic stimulation,” *Eur. J. Appl. Physiol.*, vol. 88, no. 6, pp. 601–606, Feb. 2003.
- [64] J. Domhof, J. F. Kooij, and D. M. Gavrilu, “An extrinsic calibration tool for radar, camera and lidar,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, May 2019. [Online]. Available: <https://doi.org/10.1109/icra.2019.8794186>
- [65] C.-L. Lee, Y.-H. Hsueh, C.-C. Wang, and W.-C. Lin, “Extrinsic and temporal calibration of automotive radar and 3D lidar,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 9976–9983.