Acceleration Magnification for Visualising Blood Flow Pulsation in the Skin

M.A.R.C.M. Verzijl

Technische Universiteit Delft

TU Delft

Acceleration Magnification for Visualising Blood Flow Pulsation in the Skin

by

M.A.R.C.M. Verzijl

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Thursday September 24, 2020 at 09:30 AM.

Student number:	4282604	
Project duration:	February 10, 2020 – September 24, 2020	
Thesis committee:	Prof. dr. ir. B. P. F. Lelieveldt,	TU Delft, LUMC, supervisor
	Dr. Ir. J. Dijkstra,	LUMC
	Dr. J. R. van der Vorst,	LUMC
	Dr. F. M. Vos,	TU Delft, AMC
	Dr. S. L. Pintea,	TU Delft

This thesis is confidential and cannot be made public until September 24, 2021.

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

TU Delft

Abstract

While millions of people world wide suffer from arterial diseases, such as peripheral arterial disease, there are a limited number of methods that can be used to diagnose and track these diseases which are also easy, quick and non-invasive.

This work focuses on what is needed to improve diagnosing and tracking of peripheral arterial disease (PAD) using visualisation techniques. Visualising the blood flow pulsation in the skin can be useful in cases of arterial diseases, as the diseases can influence the blood flow by obstructions and arterial stiffness. The main objective in the visualisation is to show the acceleration of the blood flow, as this is linked to the arterial stiffness.

The proposed algorithm for visualising the acceleration of the blood flow is comprised of multiple steps, including techniques such as motion reduction, Eulerian video magnification, remote photoplethysmography signal extraction and using the second derivative. The input of this algorithm are videos of the skin of patients, this makes this method easy and non-invasive.

Photoplethysmography (PPG) signals are present in videos of skin, but can not be seen with the naked eye. Using Eulerian video magnification the PPG signals are amplified for better processing and visibility. By combining groups of pixels into small patches, a decrease in processing time is achieved and it adds a filtering effect. The size of the patches controls the resolution of the visualisation. Movement from the camera or patient is detrimental when extracting the PPG signal from the video. To counteract the motion in the videos a motion reduction step, using optical flow, is applied. Using the Plane-orthogonal-to-skin (POS) algorithm, the signal extracted from videos is converted to a PPG signal. Calculating the second derivative of the PPG signal gives the acceleration of the signal. By splitting the acceleration signal into positive and negative numbers, the acceleration and deceleration of the blood flow is visualised.

Synthetic videos simulating the skin were generated in various levels of accuracy to aid the development of the algorithm and to conduct experiments. The levels range from a simple pulsating square to a moving PPG signal over a blood vessel like structure. In addition, real videos of patients were used.

The experiments show the feasibility of visualising the acceleration of the blood flow pulsation in the skin, but also highlights areas of improvements and future research. More fine-tuning of the algorithm is needed, in addition to acquiring more videos of patients with PAD before and after surgery in a controlled environment.

A working proof of concept of the algorithm is shown. It has the potential of being a novel method of diagnosing and tracking arterial diseases.

Preface

Working on this thesis has been both absolutely interesting and one of the more difficult things that I have done, making it a fitting final chapter of my years studying at the TU Delft. Even with the Corona-virus spreading across the world and making a lot of things more challenging, I look back at the last few month with a positive feeling. I am still very happy with the direction that I have chosen and I look forward to what the future will bring me.

In the past year, I have enjoyed the support of many people and I would like to mention a couple of them. First of all, I would like to thank Boudewijn Lelievedt for his supervision and pushing me to create something new. Also a big thank you to Jouke Dijkstra, for all the supportive, useful and overall fun weekly meetings over Skype. I also appreciate the work space that was provided to me at the LUMC, even though, I was not able to go there because of the Corona-virus. Thanks to Joost van der Vorst for always being enthusiast about the project and letting me attend a few surgeries. Also thanks to Marian for keeping an eye on me and helping me with implementation problems when I was stuck.

I will forever be grateful to my parents and Harlin for always being there for me and supporting me all the way through my years at the TU Delft. I have also been extremely lucky with all the emotional and practical support from Dieuwertje during my master. Finally, I would like to thank my friends and fellow students for making these years a lot more fun.

M.A.R.C.M. Verzijl
Delft, September 2020

Contents

1 Introduction

Currently there are millions of people suffering from arterial diseases, for example, peripheral arterial disease (PAD). When PAD progresses it causes severe pain and other health problems. The main principle of diagnosing peripheral arterial disease is by looking at the blood flow, as this is reduced by occlusions, but also influenced by arterial stiffness. At this moment, the methods used to diagnose PAD are limited and some are more invasive than others. Using duplex ultrasound is non-invasive and shows the flow in colour, but has a limited field of view. Contrast angiography is currently considered the golden standard, but requires the insertion of a catheter in a blood vessel to inject a contrast agent while continuously making X-ray photos. Magnetic resonance angiography is up and coming, but still requires the injection of a contrast agent via a catheter.

Videos are easy and quick to record and do not require any sort of invasive surgery. Extracting medical signal from videos is not new, but as technology improves, the research being done in this area increases too. The discovery of being able to extract photoplethysmography (PPG) signals from videos [40] has kick-started an interest in extracting more useful medical information from videos and remote PPG signals. The introduction of Eulerian video magnification [47] made it possible to see the heart beat in the skin, but is more focused towards motion.

The main objective of this thesis is to determine what is needed to improve the diagnosing of PAD, using video processing techniques such as remote PPG imaging and video magnification in combination with visualisation techniques on videos of skin. To accomplish this goal an algorithm has been designed. By taking a video of the skin of the patient and processing it with this algorithm, the acceleration of the blood flow of the skin can be made visible, which gives an indication of the stiffness of the arteries and the flow of the blood. These factors are important when looking at arterial diseases.

This thesis starts with the necessary background knowledge on peripheral arterial disease, photoplethysmography, video magnification and motion reduction in Chapter 2. In Chapter 3 the proposed algorithm for the acceleration video magnification of blood flow pulsation in the skin is explained. With the videos from Chapter 4, the experiments in Chapter 5 are conducted. Finally, this thesis is concluded in Chapter 6 with a discussion about the results, future improvement and a conclusion.

2

Background Knowledge

It is important to get a good understanding of the different aspects that belong to the topic of this thesis as it is impossible to build something new when there is no solid foundation.

The focus of this thesis, is to see if it is possible to create a new way to visualise changes in the acceleration of blood flow pulsation, especially with arterial diseases such as peripheral arterial disease. From this question multiple building blocks can be extracted that need to be researched before this question can be answered.

The questions that arise are:

- What is peripheral arterial disease and how is it currently diagnosed?
- How can the blood flow be measured through the skin?
- What is video magnification and can it be used to enhance the visualisation of the blood flow?
- How does motion estimation and reduction work?

2.1. Peripheral Arterial Disease

Peripheral Arterial Disease (PAD) can be described as obstruction of the blood flow in arteries, most common in the lower extremities. It will first be noticed by patients when, due to the obstruction of the blood flow, during movement they will experience pain and start to have a limp. More specifically, this is called intermittent claudication, which is the pain (and possible limb) that the patient is experiencing during movement of the effected extremity. In the beginning of the disease the pain will subside when the patient rests. When the PAD reaches an advanced stage the pain will not subside during rest. As the disease furthers, the continuous worsening of the obstruction of the blood flow will lead to arterial insufficiency ulcer and gangrene and ultimately amputation of limbs [28].

It has been shown by Criqui et al. that PAD has a high mortality rate in especially older men and women. While an increase in arterial diseases can be expected in older populations, it has been found that PAD significantly increases the risk of death. People who have PAD might not have any symptoms that come with the disease, however the mortality rate is still higher compared to people without PAD [8].

To check for a disease like PAD, the ankle-brachial-index (ABI) is often used. This is a relative straightforward way to check the difference in blood pressure between the arm and the ankle, as a decrease in the ankle blood pressure would indicate the existence of obstructions. By using a simple handheld Doppler ultrasound device and a blood pressure cuff, the pressure can be measured in the arms and the ankles. By dividing the pressure measured in an ankle (P_{ankle}) by the highest pressure of the arms ($P_{brachial}$) the index is derived (see eq. 2.1) [49].

$$ABI = \frac{P_{ankle}}{P_{brachial}} \tag{2.1}$$

Table 2.1 gives an overview of which values can be measured with ABI. It shows that with this technique there is a possibility to get an indication whether PAD is present in the patient. It should be

noted that there are factors that could heavily influence the results and accuracy of the ABI. Diabetic patients have a 41% higher chance of having PAD compared to non-diabetic patients [32]. However, the ABI of a diabetic patient is often falsely high compared with what is should be [2, 7, 28]. From this can be concluded that for a large percentage of patients with PAD, the ABI method of getting a diagnosis is not accurate and might give a false negative result, which subsequently leads to no further testing of PAD and letting it go untreated. It has also been found that, even though the procedure of measuring the ABI seems easy, there seems to be a lack of a proper universal protocol and lack of good education on how to perform the procedure, with a multitude of medical staff making mistakes, resulting in incorrect measurements even after extra training [49].

Table 2.1: ABI values and their interpretation with relation to PAD. [49]

ABI	Interpretation
≥1.3	Non-Compressible
1.29-1.00	Normal
0.99-0.91	Equivocal
0.90-0.41	Mild-to-moderate PAD
≤0.4	Severe PAD

There are other ways to diagnose PAD, but the biggest difference with these compared to ABI is that ABI is quick and non-invasive. Currently, the best method for diagnosing PAD is contrast angiography, it is considered to be the golden standard of PAD tests. With this method a catheter is placed in a blood vessel and while continuously making X-rays a contrast agent is delivered into the blood stream via the catheter. This is more invasive than an ABI measurement and gives the patient a relatively high radiation dose, but it could be important to get this level of accuracy [28]. A second alternative method is magnetic resonance angiography, this also uses a contrast agent, but does not use radiation and will be expected to replace contrast angiography when technology improves. The contrast agent used in this method also has fewer side-effects for people with chronic renal insufficiency [41]. Duplex ultrasound is another method that is extensively used. This technique is non-invasive and can show the flow in colour. The downside is the limited field of view. Using this technique only a small portion of an artery can be examined at once [41].

2.2. Photoplethysmography

The heart pumps blood through the body with each heartbeat. The blood vessels in the body and skin will slightly increase in size by the sudden increase of blood. This blood volume change happens all over the body and can provide interesting information about vital functions. But how can this be measured? How can a signal be extracted that shows this increase in blood volume? The answer to these questions is photoplethysmography (PPG). This section will explain what PPG is, different methods that are used to measure it and how these methods compare.

The concept of PPG was first discovered in 1937 by Hertzman, A. B. and Spealman. So while it is not a new concept in the medical world, as it has been used for decades for heart rate monitoring, the interest in PPG has increased in the last few years and the concept of PPG is being used in an increasing number of fields.

What are some of these vital functions that can be measured with PPG? First off all PPG can be used to extract import information about the heart. The first thing that is easily extracted is the heart rate (HR) and how it varies over time (Heart Rate Variability (HRV)) [31]. Even from this one simple signal extracted from a PPG signal things can be said about a patients health e.g. the functioning of the autonomic nervous system, as influences of the sympathetic and parasympathetic systems are balanced in the control of the cardiac rhythm. Irregular heartbeats could be an indication of a variety of health problems, including heart-, breathing-, kidney- and immune system problems. In the same way other variables can be extracted from a PPG signal, such as Partial Thromboplastin Time, Pulse Rate Variability and Pulse Wave Velocity. With all of these signals giving information about different aspects of the patients health [30]. Another use of PPG is that it can be used to assess the health of the arteries. For example, by using the second derivative of the PPG signal, an estimation can be made about the arterial stiffness and the acceleration of blood flow. The arterial stiffness defines the rigidity of the arteries and can give information about the health of the arteries and about various cardiovascular

diseases including PAD [30, 50].

PPG has also been used to analyse diseases by measuring blood oxygen saturation, blood perfusion, blood pressure, systolic and diastolic peaks, cardiac output, respiration, ageing, endothelial function, venous function, vasospastic conditions, vasomotor function and thermoregulation, microvascular blood flow and tissue viability and even orthostasis [1, 37]

It becomes apparent that pretty severe diseases and irregularities can be diagnosed and tracked using the signals from PPG. PPG gives the opportunity to quickly and non-invasively get information about a patient.

2.2.1. What is photoplethysmography?

After discussing the importance and applications of PPG, it is paramount to understand its working principles.

When light hits skin, some part of it will penetrate the skin and interact with the tissue and blood. Because of these interactions a part of the light will be reflected back out of the skin. Depending on the depth of the light and the amount of blood that is present in the skin there will be differences in the wavelengths that are reflected [22, 23].

There are three important parts of the skin that have an influence on the PPG signal. These three parts are the Epidermis, Dermis and Hypodermis and roughly correspond to the ranges of wavelengths for the colours blue, green and red. This means that different wavelengths of light have different penetration depths in the skin, see Figure 2.1. So at each level in the skin these light particles, with different wavelengths, interact with the tissue and blood that is present at that specific level. The light gets reflected, scattered and absorbed. For PPG the interesting part is the part that is reflected back out of the skin. When the volume of blood that flows through the vessels in the skin increases because of a heartbeat, the amount of haemoglobin also increases and this is subsequently the reason why the light that is reflected back out of the skin increases as well, because there is more blood/haemoglobin for the light to interact with and this will result in a higher amount of light being reflected. This difference in blood volume in the blood vessels can be seen when looking at the reflected light that comes back out of the skin. This reflected light can be described as a signal and is called photoplethysmography [10, 18, 22, 40, 51].

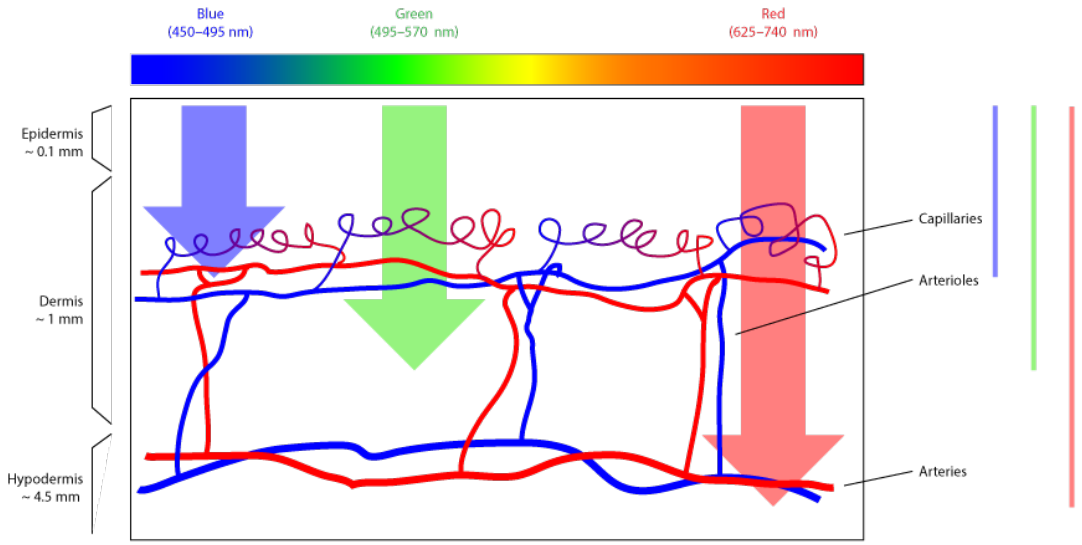

Figure 2.1: Illustration of different light wavelength and their respective depth in the skin.

2.2.2. PPG methods

The methods discussed here are the methods that follow the original way that PPG was measured, that is via a photo sensor connected to the skin. This can be done on, for example, a finger or earlobe. There are two types of physical PPG measurement techniques, the first is using transmission of light through the tissue. The light source and detector are on the opposite site of the finger/earlobe. The second way is using the reflection, in this case the light source and detector are place next to each other [37]. See Figure 2.2 for an illustration of these two types.

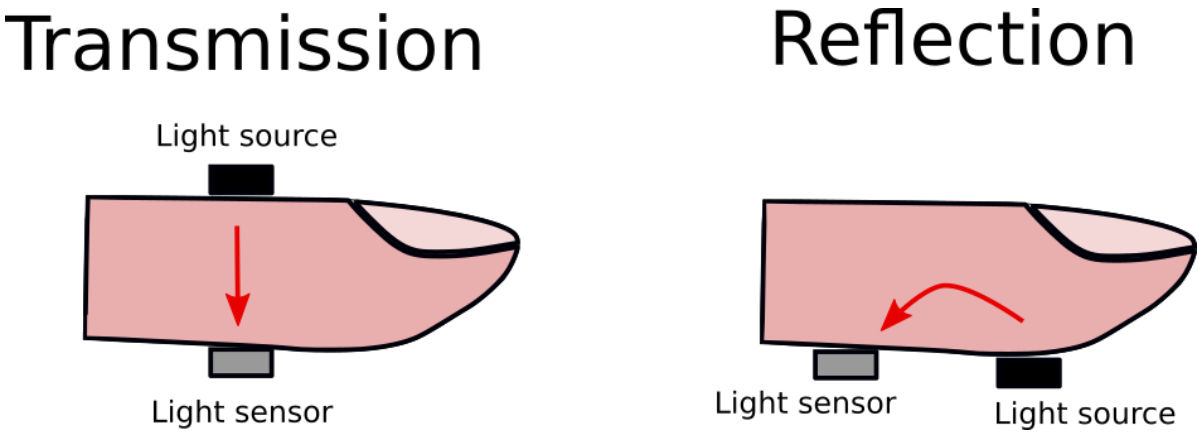

Figure 2.2: An example illustration of the two ways a physical PPG measurement can be made. The first being via transmission and the second via reflection.

2.2.3. Remote PPG imaging methods

Compared to the PPG methods that have a light sensor connected to the skin remote PPG methods have no contact with the body. Some advantages of having remote PPG methods are that simple video cameras can be used, so there is not need for special equipment. Simple webcams [34], smartphones [29] and everyday video cameras have successfully been used. Another advantage is that it is possible to get a more general PPG signal when more skin captured in a video frame compared to the local signal received from a probe connected to the skin.

There are also disadvantages, which researchers have tried to overcome or reduce. When using a dedicated light sensor that is positioned on the skin, the light that is shown on the skin is controlled by a light source that is positioned on the sensor probe. This gives the user a great amount of freedom to control the amount and wavelength of the light that hits the skin. With remote PPG there is no direct control of the light. The light that will be captured by the camera is completely depended on the light in the room [35]. If natural/ambient light is used there are many things that can influence the light. E.g. changing weather and people/objects moving around causing the illumination to change, but also certain lamps, like fluorescent tubes or dimmed LED's can cause flickering [6, 21, 38, 40]. This is often connected to the frame rate and shutter speed of the camera that is used [29]. It usually is not a problem when extracting a remote PPG signal when there is a lot of light, however, if for example, the skin is oily, wet or in some way very reflective which causes the light to immediately reflect from the skin, rather than penetrate and then reflect back out, there will be less PPG information to extract. The absence of light will also inhibit receiving a signal. When there is little to no light to interact with the skin, there will also not be any light/signal that will go from the skin to the camera.

An often overlook part of extracting a PPG signal is the contact pressure that is applied on the sensor that is placed on the skin. Grabovskis et al. shows that a non-optimal contact pressure highly influences the PPG signal and can lead to inconsistent and inaccurate results. This is another great advantage of using remote PPG methods as the pressure that is used is optimal, as there is no pressure needed during the registration of the signal [14].

The major disadvantage of remote PPG is motion. As the goal of remote PPG is to extract the minute changes in colour intensity in the skin. Any slight movement from the person in the video or from the camera will have an effect on the signal. Almost all available methods that use remote PPG have some sort of motion reduction implemented. Approaches to counteract this include taking an average over an area [46], use Region of Interest (ROI)/object tracking [10, 13, 24, 37], apply Blind Source Separation

techniques such as Principle Component Analysis or Independent Component Analysis on the PPG signal [9, 13], using different colour spaces instead of RGB [13, 45] and calculating optical flow [3, 44].

Motion is an important factor in remote PPG which cannot be ignored. There is not one golden solution to handle the motion that is present in the input videos. The research in this field is active and there are constantly new methods being developed with a specific goal and all of them are performing well in certain conditions while performing worse in other areas.

2.3. Video Magnification

Video magnification is the concept of amplifying minute changes in a video which are not visible with the naked eye. This can either be used to amplify motion or colour. It was first introduced in 2012 by Wu et al. [47]. This has sparked many implementations and variations, which have been applied in many different areas.

The following list gives the different applications of Video Magnification:

- Medical applications
  - Showing blood flow in skin and tissue. Can be used for diagnosing or following the progress of certain diseases and during surgeries such as skin replacement, but also for heart rate detection.
  - Show the breathing of babies. Can be used to check whether the baby is breathing and check the breathing rate.
- Mechanical/factory applications
  - Show the shaking of a car engine. Can be used to find any unexpected vibrations and make an engine run smoother.
- Physics applications
  - Show how vibrations move through certain materials.
  - Show how structures react to forces. For example, how a building or crane reacts to wind.
  - Basically any object that has tiny movements which cannot be physically accessed.
- Audio applications
  - Visual microphone. Extract audio from a magnified movie by looking at the changes. Sound vibrations are transmitted to a surface, these are too small/fast to see, but when magnified this can be extracted and converted to audio.

2.3.1. How does video magnification work?

In this section the general idea of how video magnification works will be explained. This follows from the original publication in 2012 [47]. The methods that followed this publication might have different ways to handle the magnification and the results might be different, but the general idea is the same. Some of these methods will be explained in the next section.

For all the video magnification methods, it is important to differentiate between the magnification of colour and motion. While most of the methods focus on accelerating the motion, some have good results in magnifying colours as well. The difficulty with amplifying the colours of the objects in the videos is that they usually still have tiny motions, which will be magnified as well.

There are four major steps that are involved in the basic video magnification. First the spatial decomposition, then the temporal processing, followed by the amplification and finally the reconstruction.

1. Spatial Decomposition In this first step an image pyramid is made. This is done by taking each frame of the input video and adding a Gaussian blur to this image and down-sampling it. By repeating this several times, depending on the desired number of levels, this creates the idea of a pyramid. In this case it is called a Gaussian pyramid. This is used to amplify the colour in a video. For amplifying the motion in a video an additional step is performed on the pyramid. By taking the difference of two

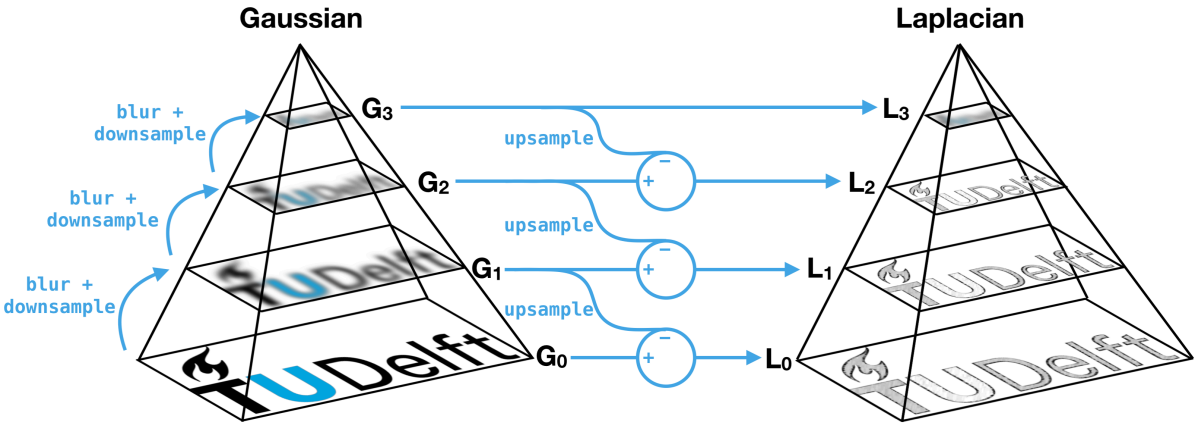

Figure 2.3: An example illustration of Gaussian and Laplacian pyramids and the steps belonging to their creation.

levels of the Gaussian pyramid and saving this in its own pyramid a Laplacian pyramid is created. See Figure 2.3.

What does this filtering and down-sampling actually do? By applying a Gaussian filter (convoluting the image with a Gaussian kernel) the image resolution is decreased. In other words this means that as a Gaussian filter is applied on an image in the spatial domain a Gaussian filter is also applied in the frequency domain. What follows from this is that the image is low pass filtered. See Figure 2.4.

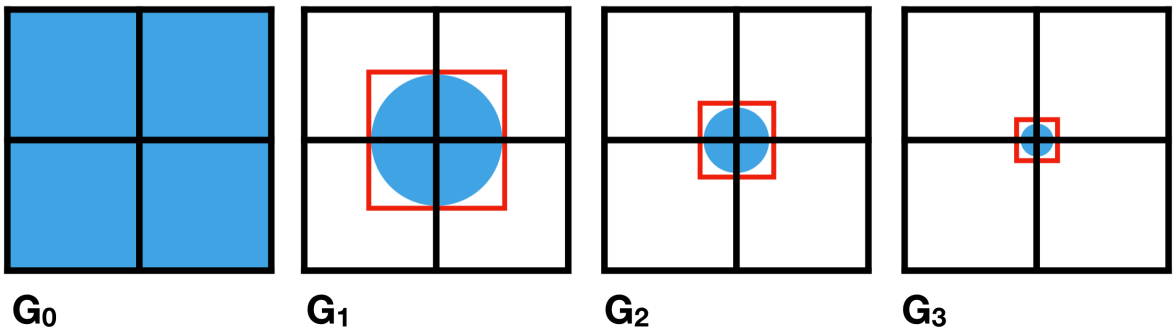

Figure 2.4: An illustration of how the sequential layers of a Gaussian pyramid in the frequency domain can look like. This shows a low pass filtering effect.

Next is how this applies to the Laplacian pyramid. As explained before each layer in the Laplacian pyramid is the difference between two layers of the Gaussian pyramid. This is done by up-sampling the previously blurred and down-sampled image of the Gaussian pyramid. See Figure 2.3. This results in each layer becoming a band pass filter, when looking at the frequency domain. Each layer of the Laplacian pyramid only contains a specific range of frequencies. This is shown in Figure 2.5.

2. Temporal processing The second major step is the temporal processing. In this step the frequency range of interest is selected. When amplifying how the blood flows through the skin, the selected frequency range will have to encompass the possible heart-rate frequencies. There are different filters that can be used, such as, ideal, IIR or Butterworth, each with a specific purpose, but the easiest, in this case, is the ideal filter. In the example of amplifying the blood flow this would mean that a frequency range from 0.8 to 1 Hz (48 to 60 bpm) should be selected to only amplify the frequencies that belong to the heart-rate, see Figure 2.6. This filter is applied to the layers of the pyramid (in the Fourier domain). Especially when using the ideal pass filtering, the outcome of this filtering is only the selected frequencies while all the other frequencies are filtered out.

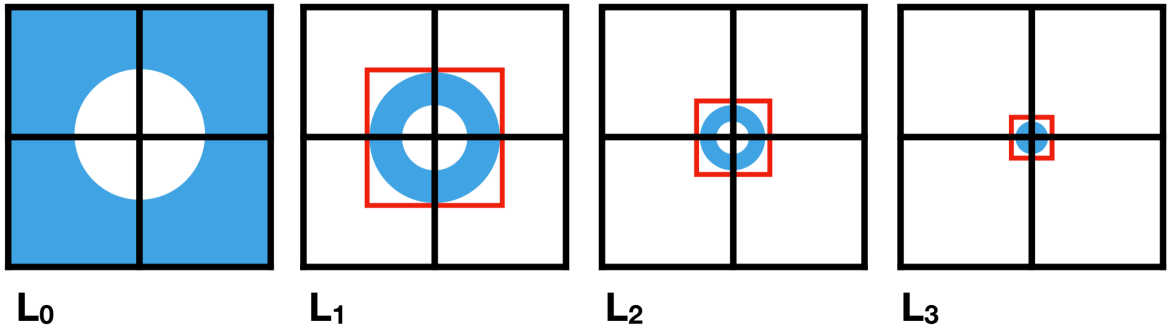

Figure 2.5: An illustration of how the sequential layers of the Laplacian pyramid in the frequency domain can look like. This shows how each layer is a band pass filter for a certain range of frequencies.

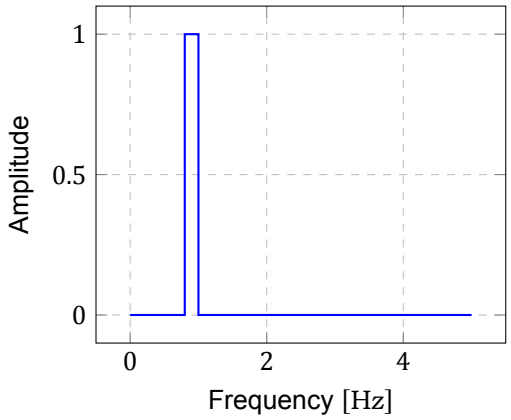

Figure 2.6: Example of an ideal filter. In this case from 0.8 to 1 Hz

3. Amplification After the layers of the pyramid are filtered the amplification step is performed. This is done by multiplying the outcome of the previous step with an amplification factor often called α. As only the interesting frequencies are present, only these frequencies will be amplified.

4. Reconstruction Before the amplification can be seen in the output video, it needs to be reconstructed. This is done by collapsing the layers of the pyramid back into a single image. This is just the opposite of before and is accomplished by up-sampling the higher layer and adding it to its lower layer, to finally arrive at the size of the original frame. After this the amplified, filtered and resized image is simply added to the original frame.

Interestingly, what is a big bottleneck in the original algorithm and other methods is the use of the Gaussian pyramids which are computational expensive and in the original implementation (as the papers suggests it should be used) uses a lot of space. However, the whole pyramid does not need to be saved (which appears to be the case when looking at the original source code and how the last level image of the pyramid is used), but instead, according to Chambino [5] a simple resize operation is sufficient. This is only the case for the colour magnification with the Gaussian pyramids, for the motion magnification the whole Laplacian pyramid still needs to be constructed.

2.3.2. What are the different techniques?

The introduction of the Eulerian Video Magnification (EVM) method [47] has sparked many improvements in this research field. Some by coming up with new methods which usually have specific focus areas, but also improvements on previous methods. As there are too many different approaches to date to be discussed in this section, a few important and/or interesting ones will be picked out and shortly explained.

Phase-Based In Section 2.3.1 the approach of Wu et al. [47] was used to explain the general workings of video magnification. A year later, from almost all of the same authors, a completely different approach was presented, called phase-based video magnification [42]. In this approach complex-valued steerable pyramids are used. Motions can be analysed by looking at the phase variations of the coefficients of these pyramids, because these variations correspond to the motion over time. After extracting these variations the steps are almost the same as in the original video magnification [47]. These are the temporal filtering, amplification and reconstruction of the video. While the name of this approach suggests it only works for amplifying motion in videos, it is mentioned in the paper that by either removing small motions with preprocessing or performing some extra steps on the amplitude, they are able to also replicate the colour magnification that was achieved by Wu et al. [47].

Building on top of the phase-based approach there are two methods that are interesting to discuss next. The first method is called Acceleration Video Magnification by Zhang et al. [52]. Zhang et al. observed that there are often large movements present which make the magnification of minute changes difficult. These large movements are usually linear, contrary to the small changes. These small changes occur on the objects that have the large linear movements. As only these small changes of the large movements need to be magnified, the acceleration of these movements is used to amplify them.

The second method, called Riesz Pyramids for Fast Phase-Based Video Magnification by Wadhwa et al. [43] was proposed to speed up the phase-based approach of Wadhwa et al. [42]. The calculations of the complex-values steerable pyramids which need to be generated for each frame, are time and resource consuming. By using Riesz pyramids instead Wadhwa et al. [43] have been able to reduce the computation time to 25%, compared to time needed by the complex-valued steerable pyramid approach. These big improvements in computation time have also made it possible for doing the computations in real time, which opened up new use cases for video magnification and makes it more attractive to be used in fast paced jobs. For example, during surgeries, training for sports or for mechanics.

Extended or changed EVM The next three methods that will be discussed are not based on the phase-based approach, but either changed or added a part to the EVM method.

The first method that changed a part in the EVM method uses Principal Component Analysis (PCA). What Wu et al. [48] did in this approach is to substitute the temporal filtering (2.3.1) with a PCA decomposition. There are two problems that this approach tried to solve. The first problem is that when tiny small changes need to be magnified with a large magnification factor, the noise that is present in the video will also be magnified significantly. The second problem that is addressed is that the user needs to manually input the frequency range of interest that will be magnified. The power of PCA is that it is able to reduce the dimensionality of the data. The way this works in the PCA-based EVM method is as follows. The spatial decomposition step removes the intensity part of the video keeping the signal that contains the changes of interest together with the noise. After performing the spatial decomposition step, to create all the different frequency bands with the residual signal, each band is passed through a PCA decomposition. By looking at all the frames at each level, the PCA is able to remove the noise dimension as well as find the dimension which has the signal of interest. After this the amplification and reconstruction steps are performed. So by the way that PCA works it automatically selects interesting signals as well as remove noise, this results in smooth magnification of tiny changes with a high amplification factor.

The second method is called Enhanced Eulerian Video Magnification [25]. In this approach EVM is used as a first step, before doing a few additional post-processing steps. After computing the magnified video with EVM, a 2D motion mapping is made by calculating the difference between the magnified and original video. This 2D motion mapping is used to create an image warping grid to amplify the motion in the video. By following the motion in the directions defined in the motion mapping, the frame pixels can be warped to amplify the motion. Using this approach the intensity of the pixels is not used for the actual magnification, but the actual movement is, which also inhibits the noise from being amplified.

The third approach of this paragraph is motion magnification using the Hermite transform [4]. Instead of the Laplacian pyramid that is used for the spatial decomposition for motion magnification in EVM, this approach uses a Hermite transform decomposition (and reconstruction). The reason why the Hermite transform is interesting is in the way that it is strongly linked to the human visual system and how a biological model of the receptive field can be described. By using this approach less noise is present in the magnified video and it is able to preserve important visual structures, such as, homogeneous region, edges and textures better than the original EVM.

Learning-Based The last approach that deserves to be mentioned is Learning-Based Video Motion Magnification [33] as it uses deep convolution neural networks. Oh et al. stated that often other approaches for motion magnification hand-designed filters are used, which are prone to noise and excessive blurring. This learning-based approach tries to learn these filters from data directly. The way that this was achieved was to carefully generate a synthetic dataset. This dataset contained source videos with very small movements and corresponding magnified videos. After putting all the videos through their deep convolution neural network to learn the filters, the filters can be used to magnify other videos, which can be real videos. Using this approach it was possible to apply motion magnification with less edge artefacts and better noise characteristics compared to other methods. One of the limitations of this approach is that when using a higher amplification factor (higher than 20) the quality of the magnification begins to degrade.

The techniques discussed above show that they are not all suitable for both motion and colour magnification. The methods available indicate a larger interest in motion magnification than colour magnification.

2.4. What is the connection between video magnification and PPG?

When putting the concepts of video magnification and PPG side by side, the connection between the two is not immediately apparent. Especially as in video magnification research the main focus is on the magnification of motion. However, when the goal of the video magnification implementation is to show the blood volume changes in the skin by magnifying the colour, the connection becomes clear. The signal that is being amplified in the colour video magnification is the PPG signal. Colour video magnification depends on the PPG signal to be present and strong enough to be picked up by the camera to get good results.

Video magnification and PPG signals have been combined in multiple ways. For example, by Sarkar

et al. [36], where they start by amplifying the video and then extracting the PPG signal. From this amplified PPG signal they do the amplification again, but this time they have a smaller frequency selection. This can be repeated to get a narrow frequency selection band which only selects the PPG signal. McDuff et al. [29] also mentions using EVM first, before extracting the PPG signal to increase the signal that is present in the video. He et al. [16] validated the use of EVM for PPG by comparing the extracted PPG signal from the video with an actual pulse sensor. All this combined gives a clear link between PPG and EVM and how they can compliment each other.

2.5. Motion in videos and how to reduce it

When recording videos of patients, which are to be used for video magnification or remote PPG, usually great care is taken to make sure the patient is sitting as still as possible and there is minimal movement from the patient (and camera). However, there will always be some movement in the patient. Even with support being present there could be breathing motion and ballistocardiographic effects (movements due to blood injected into the aorta or skin movement/deformation due to blood pulsing through larger arteries) [51]. While in some cases this motion is used to extract information about vital signs [29], in most cases, e.g. when using remote PPG, this motion will interfere with the results. Especially when information is being extracted from a single pixel or a small group of pixels, these motions will have a significant influence on the extracted signal. This means that by finding and extracting the motion of a video, it can be used to reduce or counteract the motion and thus create a clearer and more accurate signal. There are multiple techniques available to estimate the motion in videos. The important thing here is to reduce the overall/big movements while keeping the PPG information in the pixels.

2.5.1. Motion Estimation

There are two main categories of motion estimation, direct and indirect. As the name suggests, the direct methods uses the pixels of the video directly in the calculation of the motion. Whereas the indirect methods first extract features from the video and use these features to estimate the motion.

Direct Direct (or dense) methods use the variations in spatial-temporal image brightness to get the motion from each pixel. These techniques are suitable for small image motions. One technique is called block matching [20], here the current frame is divided in small blocks and in the next frame the algorithm tries to find the position where the small blocks match on this frame. The position on the new frame can then be compared to the original position of the block and this can be seen as a motion estimation. This approach is extremely computationally expensive, even with advances in search algorithm optimisation. A second approach is called phase correlation, this uses the phase information available after doing a Fourier transform. Then the cross-correlation between the two Fourier transformed images is calculated [12]. This approach works well on linear transformations but not so well on more complex motions. The third approach which is widely used is optical flow. While some call the block matching and phase correlation optical flow as well, it is usually referring to the Gunnar Farnebäck algorithm [11] which uses quadratic polynomials to approximate pixel neighbourhoods and looks how they transform or the Lucas-Kanade method [27] which assumes there is a constant movement in a local pixel neighbourhood and uses the information in these local pixels to find the motion.

Indirect Indirect methods are quite similar to direct methods, but have an added step where instead of pixels points are extracted from the frames by using techniques such as face detection, corner detection [15], SIFT detection [26]. By matching the points of two frames, the difference between these points can be calculated.

For both the direct and indirect method the difference between the positions of the pixels or points are used to build a motion mapping.

2.5.2. Motion reduction

The motion estimation methods described above all result in a motion mapping of sorts, this shows how pixels or objects are moved/transformed. This can be over time by looking at the motion between consecutive frames or by comparing all frames of the video with a single chosen frame. These motion mappings can be inverted and used to remap the frames. By doing this step the frame is transformed to the previous/reference frame and thus the motion that was present between the two frames is removed

[25]. In almost the same way ROI tracking can be used. In this case, instead of remapping the pixels to counteract the motion, the ROI is moved to align the ROIs in both frames as much as possible as to reduce the amount of motion [24].

2.6. Conclusion

During the search of the main thesis research question it became apparent that there are currently not many methods available which focus on the accelerated visualisation of blood flow pulsations in the skin, especially in combination with arterial diseases such as PAD. The aim was to determine what is needed to improve the diagnosing and tracking of PAD, by combining existing techniques into an easy to use and accessible algorithm.

Coming back to the questions that were proposed at the beginning of this chapter a foundation is created for understanding the different aspects that are involved in the main question of this thesis. PAD is a disease with a high number of patients, but it is often missed or is diagnosed wrong or too late. When the blood flow in the arteries gets restricted because of PAD this will also influence blood volume changes in the skin. Currently, the number of easy and non-invasive methods to diagnose PAD are limited.

Remote PPG methods are able to extract a signal from a video. Using a video, instead of a probe, for PPG signal extraction helps with getting a signal from a larger area and does not have the problem of making sure that the probe has the right amount of pressure on the skin. Applying video magnification can amplify the PPG signal as to make it more clear and easier to work with. The changes at pixel level are so minute, that using dense motion estimation and reduction also helps getting a more accurate PPG signal.

By using and combining the techniques found in the literature, it becomes apparent that it should be possible to show the acceleration of the blood flow pulsation in the skin and aid in diagnosing and tracking arterial diseases in an easy and non-invasive way. This could mean an addition to the currently available methods for diagnosing and tracking PAD.

3

Accelerated Flow in the Skin

The algorithm for magnifying the acceleration of the blood flow in the skin (AFitS) is made up of the following steps. The step can be divided into three categories: preprocessing, processing and visualisation.

In this section each step of the complete algorithm is explained in more detail. In the experiments chapter (5) the added value or importance of each step will be demonstrated.

- **Preprocessing**
  1. Reduce motion
  2. Select Region of Interest
- **Processing**
  3. Use EVM with a wide frequency band selection
  4. Create patches
  5. Extract rPPG signal
  6. Apply frequency filter to rPPG signal
  7. Calculate second derivative
- **Visualisation**
  8. Reconstruct

Motion Reduction Previously the influence of motion on rPPG results was discussed (see Chapter 2), this makes the first logical step (1) in this process reducing the motion in the input video as much as possible. Before any reduction to the motion can be made, the motion needs to be calculated. The motions of interest are mainly the 'larger' motions that originate, e.g. from camera movement, breathing and ballistocardiographic effects. When looking at the whole picture of what is recorded, there is usually not only skin of the patient in the frame, but also the background. This means that with the larger motions they can be found at the edges of the patient and the background. This helps to extract the correct larger motions. Combined with the main focus being on the small changes at a pixel level, optical flow was chosen for the motion estimator. If only a small patch of skin would be selected for motion reduction, it could be difficult for the optical flow to extract the larger motions. The optical flow algorithm that was chosen is from Farnebäck [11] (implemented in OpenCV). As the Farnebäck implementation of optical flow is a dense approach it looks at the pixels much more direct compared to sparse methods, this takes more time but is more accurate. This algorithm takes two frames and computes the motion between the two frames. Normally when the optical flow is calculated in a video, each frame with its subsequent are put into the algorithm, however, as the final goal of this step is to reduce the total motion of the video, each frame of the video is compared to the first frame. (This does not necessarily have to be the first frame, as long as each frame is compared to the same frame.) The

output is a motion mapping of each pixel in the x and y direction. By remapping the current frame in the opposite direction using the motion mapping the difference/motion between the two frames is removed.

The frames are converted into grey images as this is required by the optical flow implementation. At this stage there is no interest in the changes in the different colour channels, only in the motion. However, for further processing the colour channels are important. As the motion is independent of the colour channel, the remapping can be applied to the individual channels of the original colour frame to get a motion reduced colour video.

Region of interest Often not everything in a video is interesting or important to use in the algorithm, so the second step (2) is to select the region of interest (ROI). By selecting an ROI, it helps the algorithm focus on only the interesting parts and will not be distorted by things in the background. It also means that less pixels will need to be processed which will speed up the whole process. Selecting an ROI can be seen as applying a mask to a region of the frame and for each frame of the video only the area under the mask will be used. See Figure 3.1 for an example.

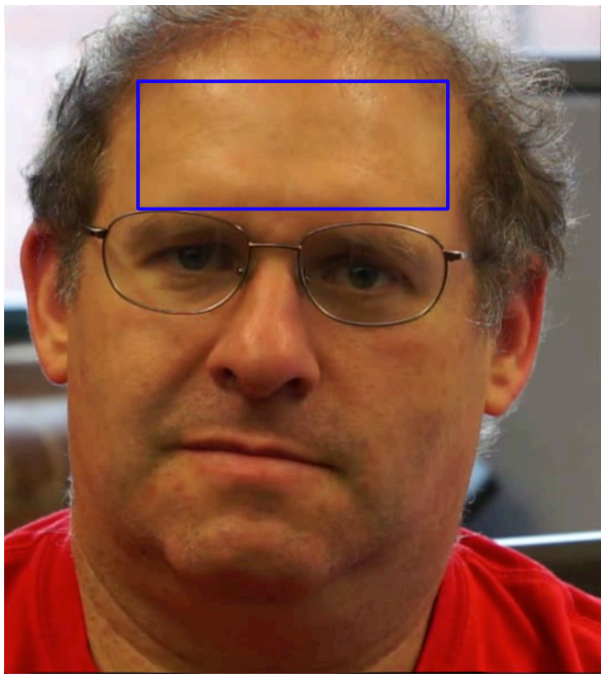

Figure 3.1: Example of a region of interest selection on a frame of a source video. Video originates from Wu et al. [47].

Eulerian Video Magnification Next are the processing steps. The PPG signal in videos can be low to start with. To make the signal stronger Eulerian video magnification (EVM) is applied (step 3). To make sure all the frequencies that could be of interest later on are magnified, a frequency range of 0.7 to 1.6 Hz is chosen.

If done correctly this will have amplified the PPG signal that should be present in the video.

Patches The next step (4) is to create patches of the video. To make sure that the resolution of the visualisation of the acceleration in the final result has a high enough resolution, it is important to take a small enough area to process at this step. By taking the whole ROI as input area in this step, it will not be possible to get the blood flow in enough detail. Doing this step pixel-wise will increase computational loads, as well as being too specific. Even with the motion reduction there will always be some motion present in the videos and this will influence signals retrieved from single pixels. To get a good amount of information, but not have small motions interfere with the signal, the best approach is to use a kernel or patch of pixels. With this approach there are multiple implementations possible with different kernels and with or without overlap, but in this case a straightforward block reduce approach was used. It takes a reduce factor and using this factor it takes a certain amount of pixels in a block and calculates the

mean of this block. The output value is a single pixel in the new output video. This step can be seen as a down sampling of the video. See step 4 in Figure 3.2. Here a patch is selected (red box) from the 4D matrix and for each frame in patch the mean is determined.

Extract rPPG signal Using the plane-orthogonal-to-skin (POS) algorithm from Wang et al. [45] in this next step (5) extracts the PPG signal in each patch over time. It combines the information from the three colour channels into one usable PPG signal. It first normalises the signal of each channel then applies the following formula, where r, g and b are the normalised red, green and blue channels respectively:

$$\begin{aligned} s1 &= g - b \\ s2 &= g + b - 2 * r \\ \alpha &= \frac{\sigma(s1)}{\sigma(s2)} \\ z &= s1 + \alpha * s2 \end{aligned} \tag{3.1}$$

The z is the raw PPG signal in a patch. See step 5 in Figure 3.2.

Filtering The output is then filtered (6) to exclude any frequencies in the signal which cannot be caused by a heartbeat. This is done by using Butterworth high and low pass filters at 0.5 and 4 Hz, this has been used by, for example, Sungjun Kwon et al. [39], Wang et al. [46] and Wu et al. [47]. To smooth the signal even further and make it more usable for the next step, a Savitzky-Golay filter is also applied. By using a window length of 19 and a polynomial order of 4 [19] the signal can be smoothed without altering the signal too much, giving the next step a more reliable signal. See step 6 in Figure 3.2.

Second derivative To extract the acceleration information that is present in the signal the second derivative of the signal is used (step 7). So when the heart beats and the blood starts to accelerate this will show a positive value and when the blood decelerates it will be negative. The OpenCV implementation of the Savitzky-Golay filter actually has a build-in derivative function. Using this function gives a better result than using a function like numpy diff. See step 7 in Figure 3.2.

Reconstruction The last step is to take the second derivative signal and reconstruct it into a layer that can be put on top of the original video. This step can be divided into a few smaller steps. As the data can contain some outliers which can influence the next step, the 1st and 99th percentiles are calculated and all values in the signal are clipped to these percentiles. Since the patched data has a smaller size compared to the original video (or if an ROI is selected, the ROI video) it first needs to be resized to make them equally sized. This can be done by up sampling the frame, so now every pixel in the patched video is mapped to a block of pixels in the original video. The values at this stage are all quite low and if these would just be added to the original video it would be hard to see, to counter this it is multiplied with an amplification factor (note: this is not the same amplification factor used in the EVM step). Next is the separation of the acceleration and deceleration component of the signal. By splitting the positive and negative values into separate channels and taking the absolute value of the negative channel, two channels are created for the acceleration and deceleration respectively. Combined with an empty channel, these channels represent the RGB channels needed. In this case, to use the red channel for the acceleration and the blue channel for deceleration was chosen arbitrarily. Finally, this layer visualising the blood flow acceleration is overlaid on top of the original video. See step 8 in Figure 3.2.

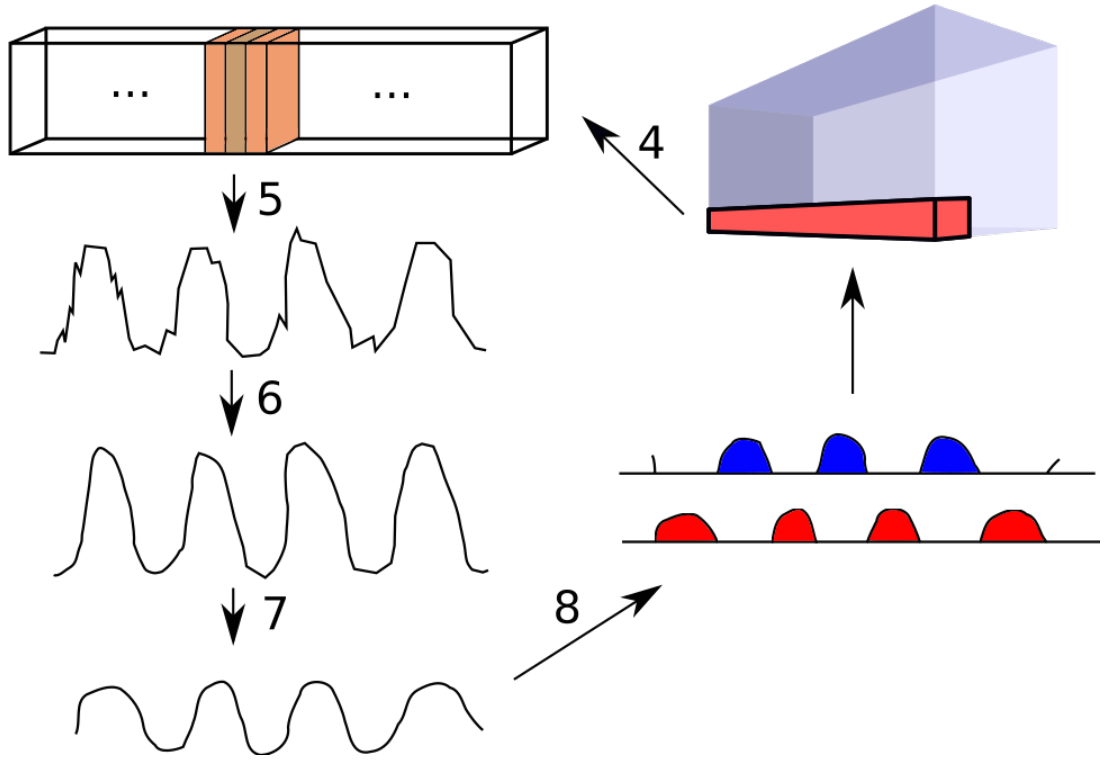

Figure 3.2: A rough and simple illustration of the steps 4-8 of the AFitS algorithm. Some small steps are not illustrated for clarity.

4 Videos

The most important thing when doing anything with remote PPG or extracting information from videos is that the videos actually contain information. If the input videos do not contain the correct or any information, even a perfect algorithm and implementation will not be able to generate the expected results.

That is why it was decided to also use synthetic videos, because all the different variables can be controlled when generating synthetic videos. This means that the ground truth of the signal is known and thus a better assessment can be made whether the results are what they are supposed to be, when looking at the results of the algorithm. It aids in proving that the concept works.

4.1. Synthetic Videos

While it is easy to find videos on the internet that could be used for video magnification of the blood flow, often there is no way of knowing whether there will be a good signal present in the videos. Another possibility is to record videos of patients, but with these videos there is no ground truth available of how the magnified version of the video should look like and how large the input signal is. There are many different video magnification implementations available and while these show some kind of magnification, there is no way to check whether this magnification is correct.

To try and combat all these problems synthetic videos can be used. Synthetic videos are videos that are generated by the computer and not filmed with a camera. This way all the variables are controlled and it is possible to generate videos with different levels of magnification.

After generating these videos, these unmagnified videos, are put through the Accelerated Flow in the Skin (AFitS) algorithm.

4.1.1. Creation process of the synthetic PPG videos

When the blood pulses through the blood vessels in the skin a PPG signal is created which can be extracted and processed. This concept of the blood pulsing in the input videos is what needs to be simulated with these synthetic videos. These videos will give the impression of the blood flowing through the skin. As the changes in the blood flow in real skin cannot be seen with the naked eye, the signal in these synthetic videos is also very low. However, the advantage of generating videos is that by increasing the magnification factor the signal can be visualised which makes checking whether the signal is correctly generated easier.

The simulated videos were made at different levels of difficulty. With difficulty meaning more accurate skin representation and real life variables. This helps to determine what the strengths and weaknesses are of the acceleration video magnification. Since any (blood flow) movement has a large impact on the magnification, the lower level videos do not have any movement, but this is added at higher levels to improve realism and test the algorithms in a more robust way.

The use of synthetic videos in research is not uncommon. However, these videos are highly specific to a certain research question. As blood flow models of the skin were not found, the videos needed to be made.

A python script was made to create these videos. There is a function for each difficulty level. In the main function the settings variables can be adjusted. These include size, fractions for the RGB channels, number of frames (video length), frame rate and amplification factor. After creating an 4D matrix consisting of all frames over time, it is saved to disk as a video file. For convenience all the videos where made 350 by 350 pixels, saved in RGB colour as this is most common in cameras and videos.

While the generated videos provide increasingly more accurate models of the blood flow through the skin, they are not accurate models of the skin and its physical properties and thus should not be used as such. These videos represent how the light is reflected out of the skin and then recorded with a video camera.

Different levels

1. A simple pulsating square where the whole square pulses in unison following a simple sine wave. This could be seen as very zoomed in part of the skin.
2. A pulsating square which also pulses in unison, but instead of a simple sine wave it uses an actual recorded PPG signal from the BIDMC PPG Dataset found on PhysioNet[1].
3. A square with the same PPG signal as in level 2, but instead of the square pulsating in unison it is a wave traversing over the square, like how blood seems to flow through the skin. This could be seen as more zoomed out version of levels 1-2 as a larger field of view is need to see the propagation of a wave.
4. Same as level 3, but in this case, a mask (or filter) is added on top of the wave which resembles the veins in the skin. This way the signal/wave looks like it is coming from the blood vessels.

As discussed in Section 2.2.1, blue light barely penetrates the skin and gets stopped in the epidermis, this is also why the fraction of blue light in the videos is lower.

Level 1 Define a frequency that can be seen as a heart rate e.g. $f = 1.33$ Hz (equal to 80 bpm) and the amplitude of the signal e.g. $a = 0.1$ as these blood volume changes in the skin are tiny.

$$\begin{aligned} y &= a * \sin(\frac{f}{2\pi}) \\ z &= y + |min(y)| + d \end{aligned} \tag{4.1}$$

Then apply Formula 4.1. Here d is an optional positive value that can be used to give an additional increase to the whole signal. The output of 4.1 (z) is an array containing the signal (only positive values), this is applied to the 4D matrix that is the output video. All the pixels in the frame that corresponds to the same index in the signal gets the signal at that index, multiplied with the colour fraction of that channel, here the background value is also added. To change the amount of amplification of the signal, an amplification factor is multiplied with the signal before applying the colour fraction and background. Figure 4.1 shows the input signal and the unamplified RGB output values.

Level 2 This level tries to simulate an actual PPG signal instead of a sine wave. By using a sample from the BIDMC dataset[1] which contains PPG signals this can be achieved. The signal is sampled at 125 Hz. To make the sample rate fit the output frame rate of the video the sample is resampled to, in this case, 30 frames per second. Then the steps are the same as in level 1. This results in a square that pulses according to the input PPG signal. See Figure 4.2.

Level 3 At level 3 the PPG signal moves along the frame over time, like a wave. Small changes are introduced to the colour channels besides the fractions.

Take the resampled PPG signal from level 2 as a starting point and define a shift factor. This shift factor is responsible for the speed of the wave procession over the frame. For each frame in the video,

[1] https://physionet.org/content/bidmc/1.0.0/

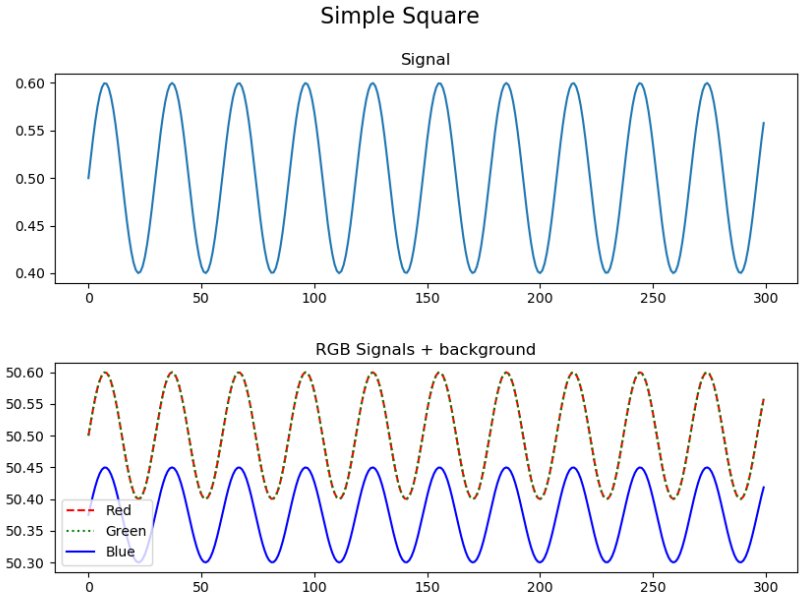

Figure 4.1: Example of the output values of level 1 over time. (In this case: $f = 80/60$, $a = 0.1$, fractions red and green are 1, blue is 0.75.)

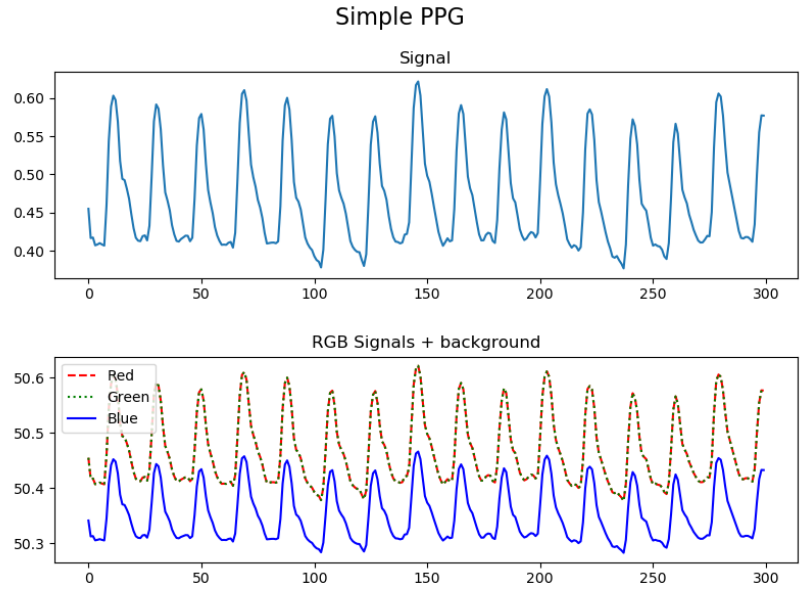

Figure 4.2: Example of the output values of level 2. (In this case: fractions red and green are 1, blue is 0.75.)

shift the signal with the shift factor times the index and from this output select a part that is the size of the width of the output video, starting at the position that is the end minus the width. For the red channel at the shifting step a red shift is added to the (shift factor times index) value. This gives the red channel an extra shift. These steps result a in 1D signal that needs to be mapped to a 2D frame. This is simply done by copying the 1D signal vertically to all the rows of the frame. Just before the frame is filled and the colour fraction is applied and the background is added, the signal is multiplied with the amplification factor. The first plot of Figure 4.3 shows the original PPG input signal, the second plot shows the amplified RGB values when looking at the first row of pixels of the first frame, this clearly shows the wave. The third plot shows the amplified values of the middle pixel of the first row over time.

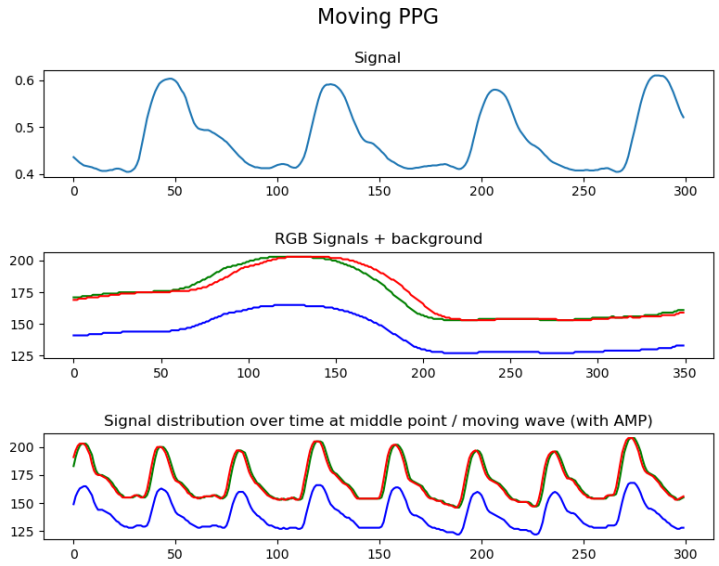

Figure 4.3: Example of the output values of level 3. (In this case: fractions red and green are 1, blue is 0.75, shift factor is 10 and red shift is also 10.)

Level 4 For the last level, the only addition compared to level 3, is that a mask in the shape of a vascular system is added to the signal. An image of blood vessels [2] was processed with a photo editor to create a mask image. Figure 4.4 shows what an amplified frame looks like with the mask in place. The first plot of Figure 4.5 shows a sample from the input signal matched to the scale of the output signal. The second plot shows amplified values of the first frame over the width, the mean is taken in the vertical direction because of the mask. Without the mean there would be drops to zero in the signal. Even with the mean there is noise in the signal as the mask is not uniform when comparing vertical slices. The third plot shows the vertically averaged amplified values of the middle column over time.

[2] https://www.asianscientist.com/2019/02/in-the-lab/nanoparticles-vascular-leak-cancer-tumor-metastasis/

Figure 4.4: Example frame of a level 4 video. (In this case: fractions red and green are 1, blue is 0.75, shift factor is 10 and red shift is also 10.)

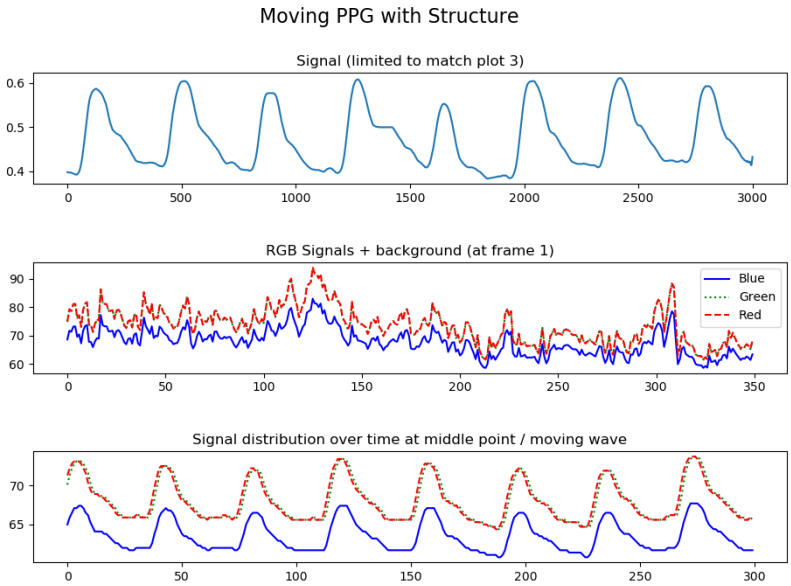

Figure 4.5: Example of the output values of level 4 at different positions and timestamps. (In this case: fractions red and green are 1, blue is 0.75, shift factor is 10 and red shift is also 10.)

Future improvements To improve the accuracy of the blood flow simulation, the next step would be to differentiate between the spread of the different colour channels through the skin. Depending on the colour channel the mask should spread out more over the frame. Red light reflects deeper and thus gives information about the bigger arteries. To simulate this the signal in the red channel should follow the blood vessel mask more closely. When the depth of the reflected light decreases (with respectively green and blue light), the 'blood' should flow more with a blurred border along the blood vessel mask, simulating the blood flowing through the arterioles and capillaries.

While it might seem intuitive that the blood flows in one direction in the blood vessels this is not really the case when looking at the skin as a whole. When looking at the skin (Figure 2.1) the arterioles and especially the capillaries are more randomly positioned and do not necessarily follow the main blood flow of the larger arteries. In the simulation, this would result in a stronger directional flow in the blood vessel mask but a slightly delayed and more random flow in the smaller blood vessels.

However, creating these levels of skin simulation is outside of the scope of this thesis and could be the topic of a thesis by itself.

4.1.2. Creation process of synthetic videos with small motion

To test the motion reduction step of the algorithm synthetic videos of a block with motion were generated. These videos were used to check whether the motion reduction actually worked as expected and to find the limits of the motion reduction implementation.

To create the effect of a moving square a sine wave signal is generated and this signal is used to change the position of the origin of the square in the x and y direction. By adding a factor to the y direction, the proportion between horizontal and vertical movement of the square can be altered.

The colours of both the background as well as the moving square itself are also adjustable. This can be used to see whether this influences the motion reduction, especially when the edge between the moving square and the background gets less prominent. See Figure 4.6 for a few examples.

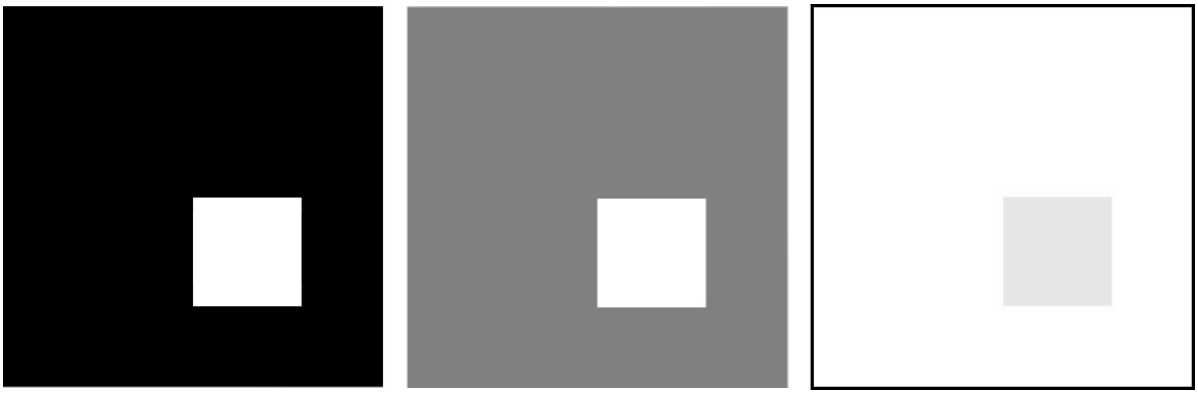

Figure 4.6: Example stills of three different synthetic videos of a moving square. The first with a black background and a white square, the second has a grey background and a white square and the last example has a white background with a grey square.

4.2. Other Videos

To test how the AFitS algorithm performs in real world applications, real videos were selected. The first real videos that were selected originate from the original EVM paper [47]. The videos that were used in this paper can be found online on the website about this method[3]. See Figure 4.7 for screenshots of two examples of these videos.

From the Leiden University Medical Center (LUMC) multiple video dataset were acquired. One dataset contains videos of the feet of patients before and after undergoing thromboendarterectomy (TEA). TEA is an operation where clotted blood is removed from arteries. Another dataset contains videos which can be used as a control, these contain mostly feet and a few hands.

[3] https://people.csail.mit.edu/mrub/evm/

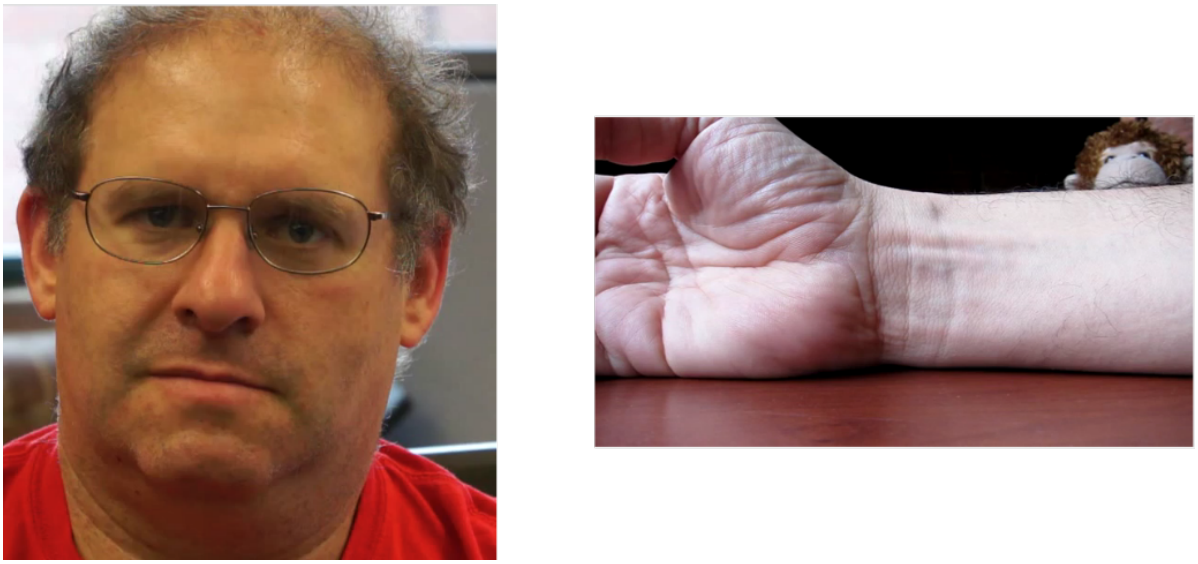

Figure 4.7: Example stills of two videos from the original EVM paper [47].

5

Experiments

The experiments using the algorithm and videos are divided into multiple parts. In the first part synthetic videos are used. For the second part real videos with skin will be used. The third part will be about a video of a patient before and after a TEA. In the fourth part individual steps of the algorithm will be bypassed to see what the influence of these steps on the results are.

5.1. Part 1

In this part the synthetic videos are used. The main focus will be on the plots that were generated for each step of the algorithm.

From the four different levels of skin simulation mentioned in Section 4, only the last two levels will be discussed in detail. Due to the frames of the videos pulsating in the same colour, the results from the algorithm for the first two levels are difficult to portray in a report. The third and fourth level offer more insight as blood flow simulation is more accurate.

Figure 5.1 shows screenshots of frame from the third and fourth level synthetic videos. As the PPG signal is extremely low and the background is a uniform colour the images appear to be almost completely dark grey. This shows the feasibility of generating videos with a minute signal which cannot be seen with the naked eye. The background colour is not the focus of these synthetic videos as it should not have an influence on the PPG signal itself. Having this particular background colour does aid in seeing the visualisation effect in the output of the algorithm more clearly.

(a) Level 3

(b) Level 4

Figure 5.1: Screenshots of input frames from the third and fourth level synthetic videos.

In the case of these synthetic videos there is no motion present in the form of patient or camera movement. Thus the motion reduction step was skipped for both videos.

In Figure 5.2 a difference can be seen between the signals from the two levels. In the top plot of 5.2b the signal seams to be none existent. This is correct, as in this level a mask is used to simulate the blood vessel structure and when the middle pixel was selected to show the signal, a pixel was chosen which was not behind the mask and thus has no signal. To handle this issue the mean over the y-axis is calculated too. The signal does not seem very clear and smooth, this is because the changes of intensity are so minute and in the way that video files are saved, there are often only 256 different levels of intensity per channel.

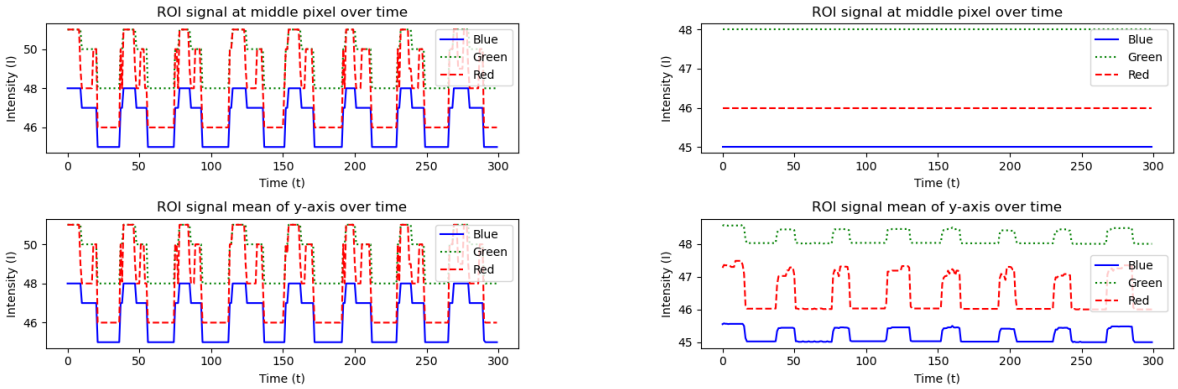

(a) Level 3 (b) Level 4

Figure 5.2: Results of the ROI selection of the third and fourth level synthetic videos.

By applying the video magnification step it becomes apparent how much the signal improves. Eulerian video magnification is able to select the interesting frequencies and amplify them. This also clearly shows the extra filters that are used in video magnification as the signal is much smoother.

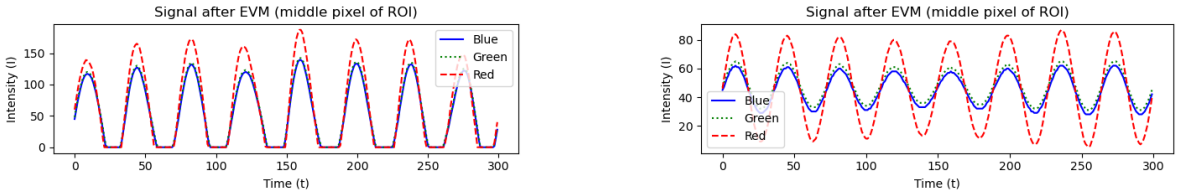

(a) Level 3 (b) Level 4

Figure 5.3: Results of the video magnification step of the third and fourth level synthetic videos.

In the case of these synthetic videos the patching does not seem to alter the signal at all. In a (small) neighbourhood the pixels are almost all the same value, so the mean value of this group will be very close to the value of each of the pixels. Also there is no noise in these synthetic videos that can be filtered out with the averaging.

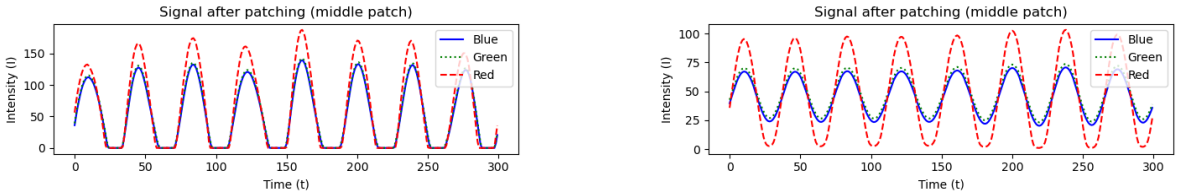

(a) Level 3 (b) Level 4

Figure 5.4: Results of the patching step of the third and fourth level synthetic videos.

In the last figure displaying plots for the synthetic videos (Figure 5.5) there are unexpected differ-

ences in the signals, especially in the frequency of the POS signals. Compared to the real videos, discussed next, this mostly appears to be happening with the synthetic videos. This points to a limitation in the current implementation of the synthetic videos. However, when looking at the resulting videos as a whole, instead of at a single point, the results seem promising. In Figure 5.6 the red acceleration and blue deceleration is clearly visible and in 5.6b the resolution due to the patches and the accuracy due to the mask is clearly visible, giving a good indication that this algorithm show the actual accelerated blood flow pulsation in the skin.

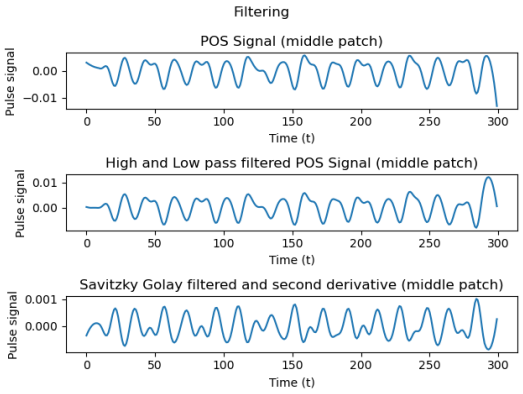

(a) Level 3

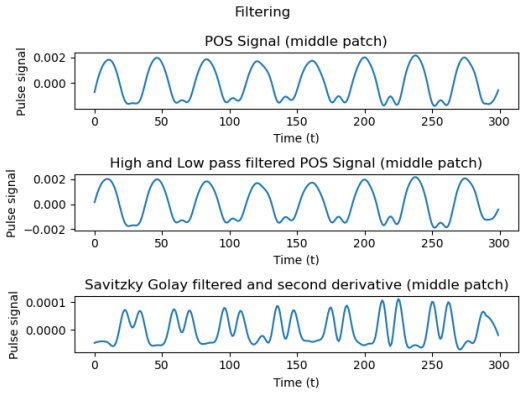

(b) Level 4

Figure 5.5: Results of the filtering and second derivative steps of the third and fourth level synthetic videos.

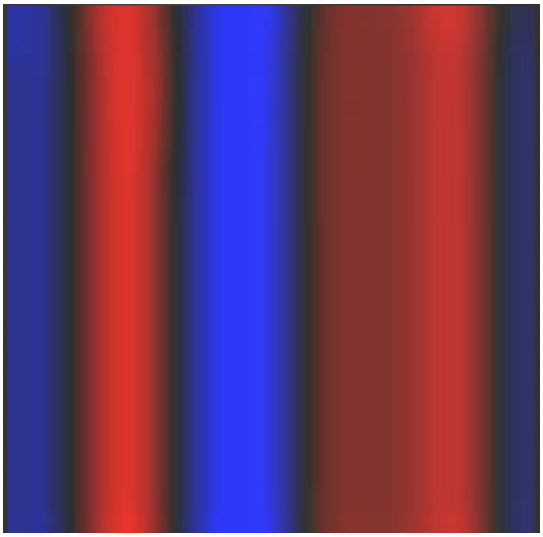

(a) Level 3

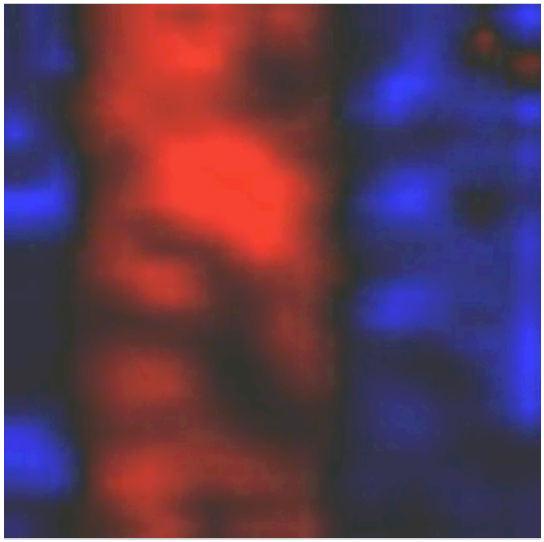

(b) Level 4

Figure 5.6: Screenshots of the final result of the third and fourth level synthetic videos.

5.2. Part 2

For the experiments with the real videos, the face and wrist videos from Wu et al. [47] were chosen. Motion reduction was applied to both videos.

In Figure 5.7 the ROIs that were selected are shown. These ROIs were chosen to get a clear area of skin with minimal background or other objects.

When looking at the signals in Figure 5.8 it seems like there is almost no signal present. It is important to notice the difference of scaling in the colour channels between 5.8 and 5.2. In the synthetic videos the amplitudes of the channels are relatively close to each other, in the real videos the amplitudes have a larger difference. However, the intensity changes, in this case, are still minute.

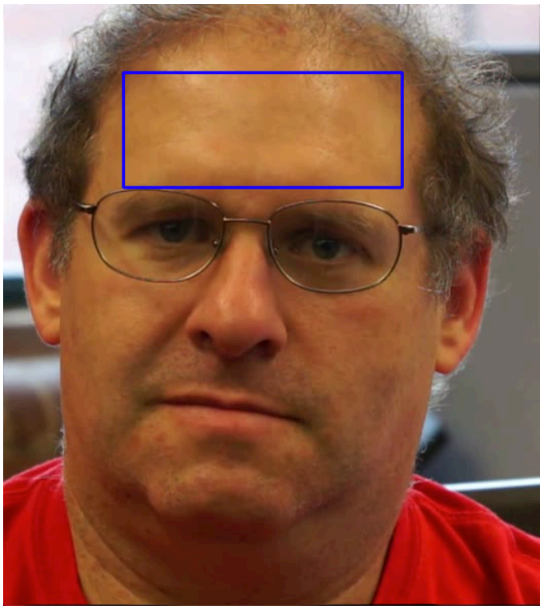

(a) Face

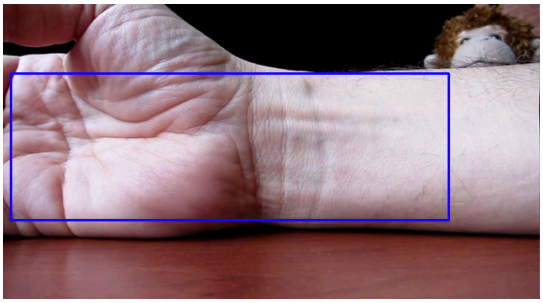

(b) Wrist

Figure 5.7: Screenshots of input frames from the face and wrist videos.

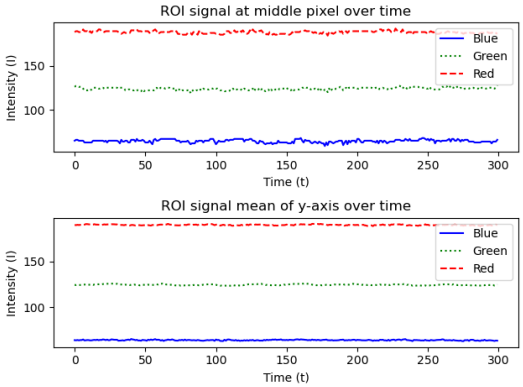

(a) Face

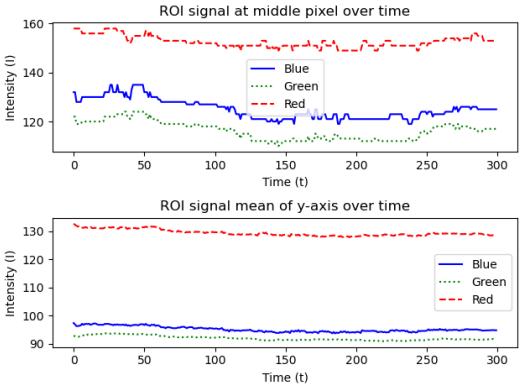

(b) Wrist

Figure 5.8: Results of the ROI selection of the face and wrist videos.

In the real videos the video magnification shows a significant improvement in signal. Looking at the relatively smooth signal the assumption can be made that there is a usable PPG signal present in these videos.

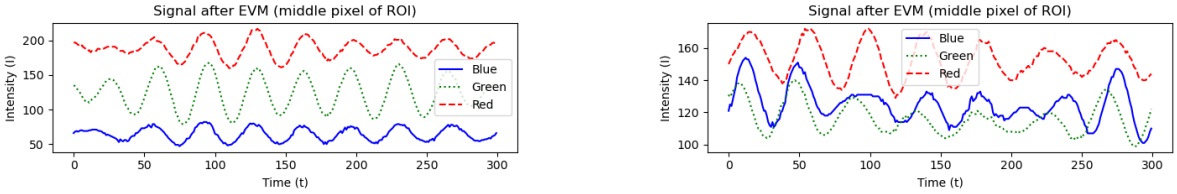

(a) Face (b) Wrist

Figure 5.9: Results of the video magnification step of the face and wrist videos.

Where with the synthetic videos the patching step did not seem to be doing much, in the real videos this step does seem to smoothen the signal by getting rid of high frequency noise.

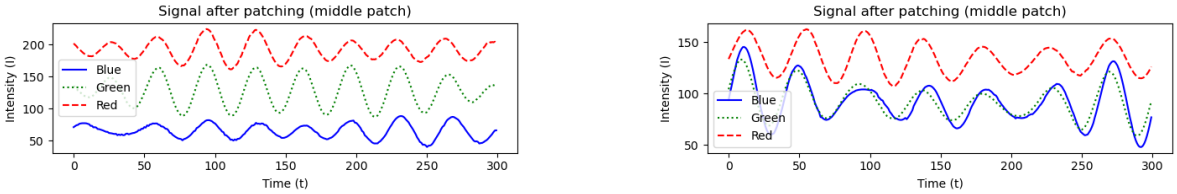

(a) Face (b) Wrist

Figure 5.10: Results of the patching step of the face and wrist videos.

Apart from some small artefacts that can be found in the plots of Figure 5.11, the signals look much more realistic and usable than those of the synthetic videos (Figure 5.5).

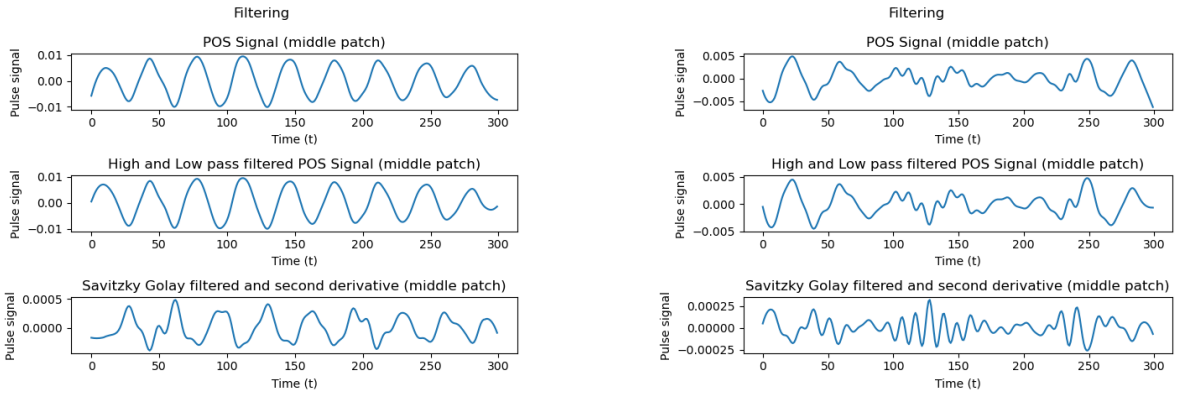

(a) Face (b) Wrist

Figure 5.11: Results of the filtering and second derivative step of the face and wrist videos.

The results clearly show the amplification. Intriguingly, the pulses can be seen flowing through the skin. Immediately, the acceleration and deceleration blood flow patterns are noticeable.

(a) Face

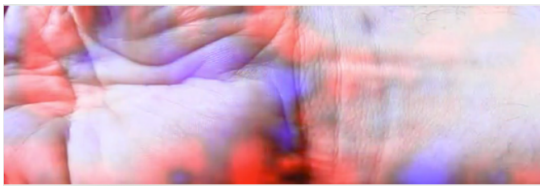

(b) Wrist

Figure 5.12: Screenshots of the final result of the face and wrist videos.

5.3. Part 3

The focus of this part is on visualising the difference between videos of a patient with PAD before and after undergoing a TEA treatment.

(a) Before TEA

(b) After TEA

Figure 5.13: Screenshots of input frames from the before and after TEA videos.

After selecting an ROI (Figure 5.13) and applying the video magnification and the patching steps, Figure 5.14 shows similar signals for both before and after the TEA (ROI and EVM results left out, as these are similar to plots of previous experiments). Most noticeable is the larger difference between the red channel compared to the blue and green channels. From the plots in Figure 5.15 it is hard to tell if, at these specific locations, there are actually increases in signal. The stills in Figure 5.16 also show no significant difference. Looking at the actual videos it is equally difficult to pick out differences that could be the result of the TEA.

In Section 6 possible reasons why in these two videos the difference does not seem so noticeable will be discussed more in-depth.

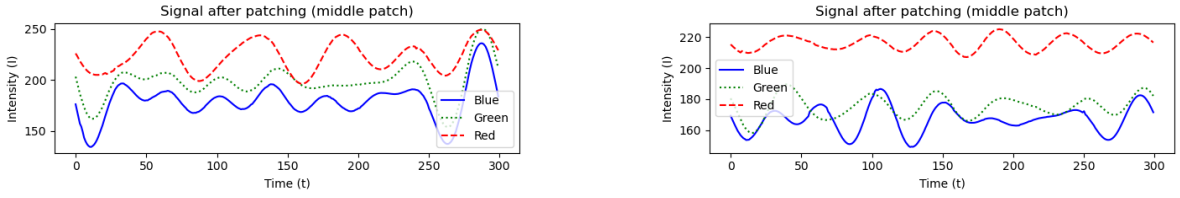

(a) Before TEA

(b) After TEA

Figure 5.14: Results from the patching step of the before and after TEA videos.

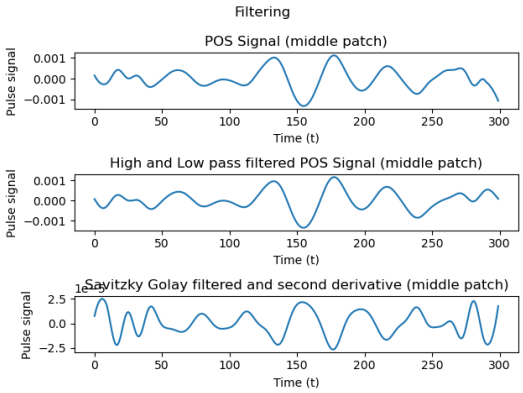

(a) Before TEA

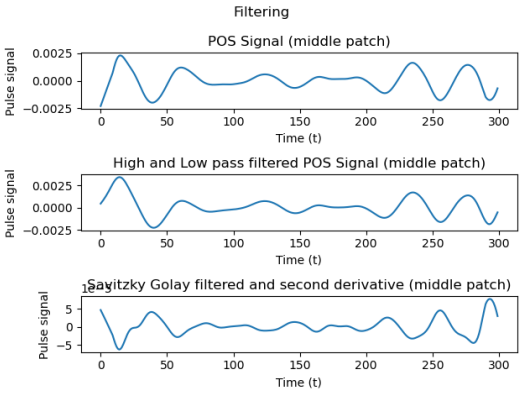

(b) After TEA

Figure 5.15: Results from the filtering and second derivative step of the before and after TEA videos.

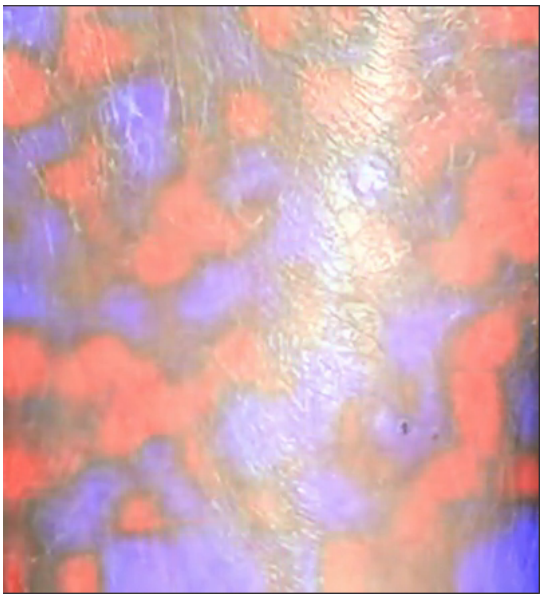

(a) Before TEA

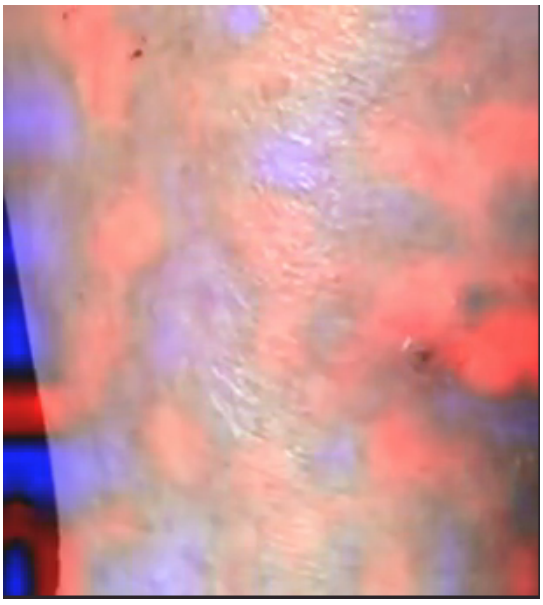

(b) After TEA

Figure 5.16: Screenshots of the final result of the before and after TEA videos.

5.4. Part 4

For the last part of the experiments certain steps in the algorithm are bypassed.

Without motion reduction Most of the videos in the available collection have very little motion in them. For this experiment a video was selected which contained the most visible amount of motion (See Figure 5.17).

From the filtered plots in Figure 5.18 there does not seem to be a significant difference between the signals of the same video, but one with and one without motion reduction. The difference in signal can most likely be contributed to marginally different ROIs and thus also a slightly different patches. As the amount of motion in the video without the motion reduction is quite low, it is possible that the patching has also filtered out most of the motion (See Figure 5.19).

The resulting videos do not seem to be affected greatly by the motion. It is noteworthy, that using the motion reduction does give a better experience when watching the videos, the motion can be distracting from the flow visualisation.

Figure 5.17: Screenshot of an ROI selection of a video containing visible motion.

Without video magnification The previous experiments have revealed the significant improvements the video magnification makes to the signal. By bypassing the video magnification step the results will inherently differ. Figure 5.20 shows this clearly. The signal has more noise and is slightly less smooth. The difference is also noticeable in the output videos. The flow visualisation in the video without the video magnification flows much faster, more erratic and with small blobs, making the patterns in the visualisation difficult to see or follow.

Without patching Bypassing the patching step is immediately noticeable due to the extreme increase in computational time for each of the following steps in the algorithm. Except for the motion reduction step, which processes the whole video frame, performing each step normally takes a maximum of a

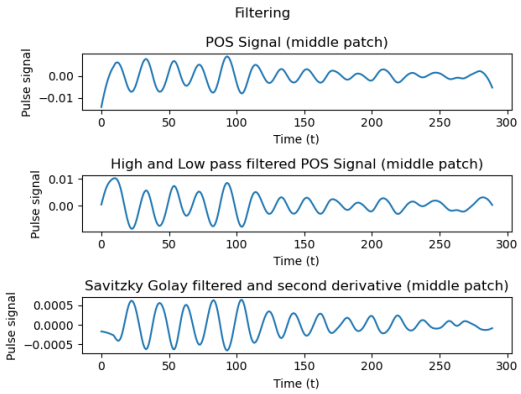

(a) With motion reduction.

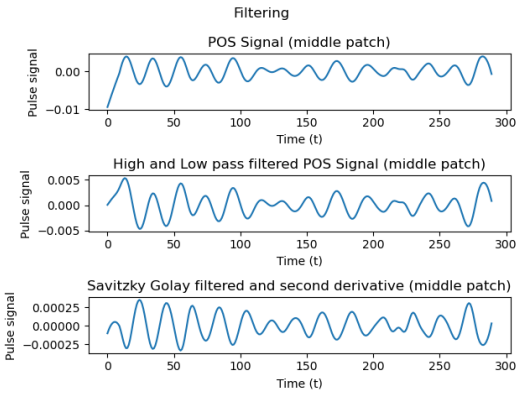

(b) Without motion reduction.

Figure 5.18: Results of the filtering and second derivative steps without and with motion reduction.

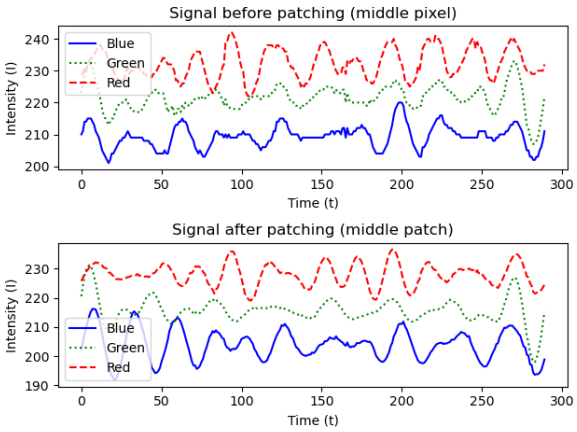

Figure 5.19: Results of the patching step of a video containing visible motion.

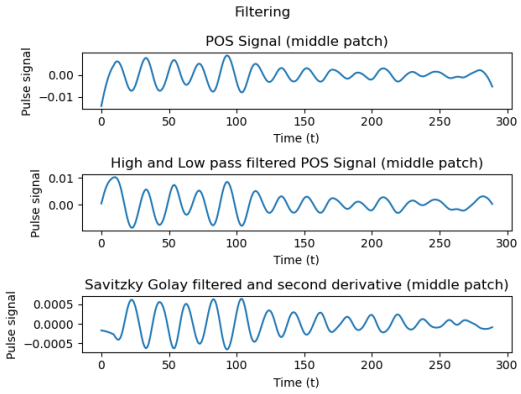

(a) With video magnification.

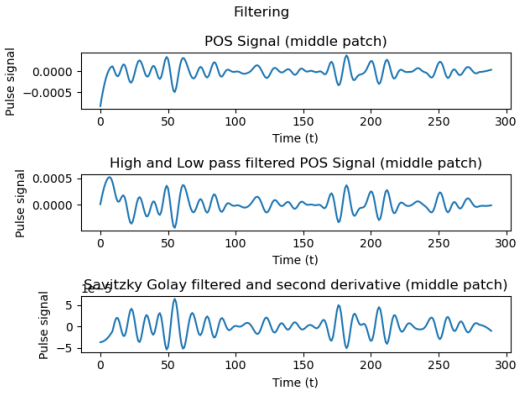

(b) Without video magnification.

Figure 5.20: Results of the filtering and second derivative steps without and with motion reduction.

few seconds. This changes drastically when not creating patches, especially the filtering and second derivative steps can take minutes even with a small ROI. It might even seem like the program has come to a stop. In Figure 5.21 artefacts can be seen which are not present in the patched version, this can be attributed to the additional filtering the patching provides. An advantage of processing all the pixels individually is that the resolution of the visualisation increases. While videos like 5.17 can benefit from this, other videos like the face video from Wu et al. [47] do not need the extra resolution. It can even be more difficult to see the flow and follow it easily. When applying this to a synthetic video (moving PPG with structure) the resolution is much higher, but a slight distortion is introduced (see Figure 5.22).

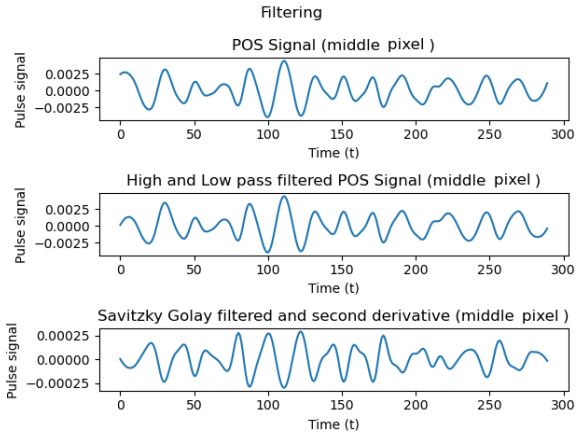

Figure 5.21: Results of the filtering and second derivative step of a video without patching.

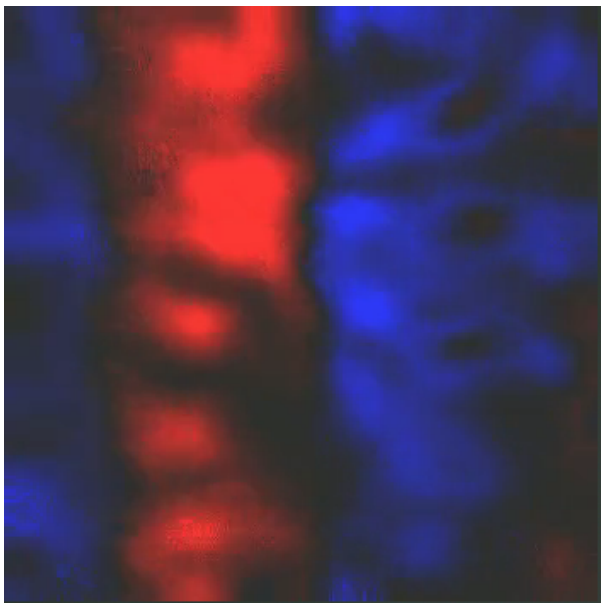

Figure 5.22: Screenshot of the output of a synthetic video without patching.

5.5. Conclusion experiments

The experiments described above show interesting results when using this algorithm to magnify the acceleration of the blood flow pulsation in the skin, but they also highlight some of its pitfalls.

By using synthetic videos the concept has been demonstrated to work. With the current implementation and available videos it is difficult to see differences in acceleration of the blood flow pulsation in the skin in videos of a patient before and after a TEA. Bypassing certain steps in the algorithm has significant effects on the resulting videos. The extent of the changes due to the bypassing of certain

steps is also heavily influenced by the input videos.

6

Conclusion

Discussion and future improvements While the proof of concept has been shown to work, there are variables in the algorithm that can be changed and improved, but implementing those and fine-tuning them is outside the scope of this thesis. In the implementation, functions are used from popular libraries (e.g. Scipy and Numpy). The input parameters of these functions can be fine-tuned and could result in a more accurate visualisation. Examples of parameters which can be fine-tuned are: video magnification amplification factor, video magnification frequency range, reconstruction magnification factor, optical flow settings, Savitzk-Golay filter window length and polynomial order and patch size.

The current implementation of the algorithm was made to proof the concept of this thesis as well as perform the experiments. Even though the focus was not to optimise for speed, the program can perform all the steps, except the motion reduction, in seconds. This means that with speed optimisations it can definitely be used in (almost) real-time applications.

To save time and computational power, the current implementation limits the videos to 300 frames. The main reason for this, is that longer videos can consume too much memory and stop your computer from working. One way to counteract this would be to use a moving window, which only processes a section of the video and when done, saves that section to the disk and continues with the next section. When longer videos are being processed there is a possibility that the PPG signal increases or decreases over time. A solution to removing trends in the signal is applying a detrending algorithm, as a signal with trends in it could possibly distort the visualisation of the acceleration.

The importance of the camera used to capture the videos is something that should not be underestimated. Even though, cheap cameras, phones and webcams have successfully been used in the past, there are cameras that have a hard time capturing useful videos. Experiments have been done by looking at the signals extracted from videos captured using GoPros, but apparently these cameras (especially the newer models) do some processing on the video before being saved, which destroys any valuable PPG information. The same principle applies to certain (newer) phone models. While most capturing devices have the option to control video settings, these settings did not grant the ability to disable the processing done to the video before saving it on the device. Different compression techniques for videos also come into play here, as different compression techniques have varying levels of influence on the signal. While the choice of camera should be kept in mind, there are plenty of cheap and older cameras available, which will record the video without extreme processing. This helps making this method even more easy to use and accessible, as no expensive camera should be needed. Using a camera that supports capturing in a raw video format or using an uncompressed live feed directly streamed from the camera sensor, gives the highest probability of obtaining a clear and usable signal.

From the experiments conducted on the videos of the patient before and after the TEA, it is apparent that it is currently difficult to quantify the results in such a way that it can be used for diagnosing PAD and treatment tracking. This calls for more research where, for example, the other foot is used as a reference point, or by having a more fixed light, camera and support setup as to reduce the number of variables that can interfere with the video capture. To quantify the results it is of the upmost importance to have a point of reference. This should be used to calibrate the recordings from, for example, before and after surgery or when following the progress of a disease over time, but also when comparing

results between different patients. The access to videos from both before and after arterial disease surgeries at the moment of publication is limited. The videos used in this experiment are of the same patient, but most of the variables discussed above are not fixed between the two videos. Variables including the position of the foot relative to the camera and the lighting are different in both videos. It is possible that the camera used to record these videos is not suitable enough for this purpose.

Conclusion Despite the above-mentioned points about the limitations and possible improvements to the method, it has been proven that the concept to visualise the acceleration of blood flow pulsation in the skin is possible. By using synthetically generated videos in combination with real (patient) videos a quick and straightforward algorithm has been created. While quantifying the result at this stage has proven to be difficult, the quality of the results shows great potential for diagnosing and tracking arterial diseases. In the future, after fine-tuning the algorithm and making it more robust, the goal is to make this novel technique available to the world to aid in reducing the pain and suffering caused by PAD and other arterial diseases.

A

Implementation Details

For this thesis two python programs were made, one for generating the synthetic videos and one for the algorithm. Anaconda[1] was used for the package management. The source code of both programs (on GitHub: `https://github.com/Verzijl`) and this thesis will be made available online after a few months embargo due to another paper that needs to be published.

All the generating of the videos and the experiments of this thesis were done using a late-2013 15" Retina MacBook Pro running macOS Catalina 10.15.5.

[1] `https://anaconda.org/`

Bibliography

[1] John Allen. Photoplethysmography and its application in clinical physiological measurement. *Physiological Measurement*, 28(3):R1–R39, mar 2007. ISSN 0967-3334. doi: 10.1088/0967-3334/28/3/R01. URL https://iopscience.iop.org/article/10.1088/0967-3334/28/3/R01.

[2] John Allen, Klaus Overbeck, Alexander F. Nath, Alan Murray, and Gerard Stansby. A prospective comparison of bilateral photoplethysmography versus the ankle-brachial pressure index for detecting and quantifying lower limb peripheral arterial disease. *Journal of Vascular Surgery*, 47(4):794–802, apr 2008. ISSN 07415214. doi: 10.1016/j.jvs.2007.11.057. URL https://linkinghub.elsevier.com/retrieve/pii/S0741521407019325.

[3] Fernando Andreotti, Alexander Trumpp, Hagen Malberg, and Sebastian Zaunseder. Improved heart rate detection for camera-based photoplethysmography by means of Kalman filtering. *2015 IEEE 35th International Conference on Electronics and Nanotechnology, ELNANO 2015 - Conference Proceedings*, pages 428–433, 2015. doi: 10.1109/ELNANO.2015.7146951.

[4] Jorge Brieva, Ernesto Moya-Albor, Sandra L. Gomez-Coronel, Boris Escalante-Ramírez, Hiram Ponce, and Juan I. Mora Esquivel. Motion magnification using the Hermite transform. *11th International Symposium on Medical Information Processing and Analysis*, 9681(December 2015): 96810Q, 2015. ISSN 1996756X. doi: 10.1117/12.2209199.

[5] Pedro Boloto Chambino. *Android-based implementation of Eulerian Video Magnification for vital signs monitoring*. PhD thesis, Universidade do Porto, 2013. URL http://p.chambino.com/dissertation/pulse.pdf.

[6] Xun Chen, Juan Cheng, Rencheng Song, Yu Liu, Rabab Ward, and Z. Jane Wang. Video-Based Heart Rate Measurement: Recent Advances and Future Prospects. *IEEE Transactions on Instrumentation and Measurement*, 68(10):3600–3615, oct 2019. ISSN 0018-9456. doi: 10.1109/TIM.2018.2879706. URL https://ieeexplore.ieee.org/document/8552414/.

[7] Cécile Clairotte, Sylvie Retout, Louis Potier, Ronan Roussel, and Brigitte Escoubet. Automated Ankle-Brachial Pressure Index Measurement by Clinical Staff for Peripheral Arterial Disease Diagnosis in Nondiabetic and Diabetic Patients. *Diabetes Care*, 32(7):1231–1236, jul 2009. ISSN 0149-5992. doi: 10.2337/dc08-2230. URL http://care.diabetesjournals.org/cgi/doi/10.2337/dc08-2230.

[8] Michael H. Criqui, Robert D. Langer, Arnost Fronek, Heather S. Feigelson, Melville R. Klauber, Theresa J. McCann, and Deirdre Browner. Mortality over a Period of 10 Years in Patients with Peripheral Arterial Disease. *New England Journal of Medicine*, 326(6):381–386, feb 1992. ISSN 0028-4793. doi: 10.1056/NEJM199202063260605. URL http://www.nejm.org/doi/abs/10.1056/NEJM199202063260605.

[9] G de Haan and A van Leest. Improved motion robustness of remote-PPG by using the blood volume pulse signature. *Physiological Measurement*, 35(9):1913–1926, sep 2014. ISSN 0967-3334. doi: 10.1088/0967-3334/35/9/1913. URL http://stacks.iop.org/0967-3334/35/i=9/a=1913?key=crossref.f46064b855e94d2d426c3c66e22da0ba.

[10] Qiang Fan and Kaiyang Li. Non-contact remote estimation of cardiovascular parameters. *Biomedical Signal Processing and Control*, 40:192–203, 2018. ISSN 17468108. doi: 10.1016/j.bspc.2017.09.022. URL https://doi.org/10.1016/j.bspc.2017.09.022.

[11] Gunnar Farnebäck. Two-Frame Motion Estimation Based on Polynomial Expansion. In *Lecture Notes in Computer Science*, volume 2749, pages 363–370. 2003. ISBN 978-3-540-40601-3. doi: 10.1007/3-540-45103-X_50. URL http://link.springer.com/10.1007/3-540-45103-X{_}50.

[12] Hassan Foroosh, J.B. Zerubia, and Marc Berthod. Extension of phase correlation to subpixel registration. *IEEE Transactions on Image Processing*, 11(3):188–200, mar 2002. ISSN 10577149. doi: 10.1109/83.988953. URL http://ieeexplore.ieee.org/document/988953/.

[13] R. M. Fouad, Osama A. Omer, and Moustafa H. Aly. Optimizing Remote Photoplethysmography Using Adaptive Skin Segmentation for Real-Time Heart Rate Monitoring. *IEEE Access*, 7:76513–76528, 2019. ISSN 21693536. doi: 10.1109/ACCESS.2019.2922304.

[14] Andris Grabovskis, Zbignevs Marcinkevics, Uldis Rubins, and Edgars Kviesis-Kipge. Effect of probe contact pressure on the photoplethysmographic assessment of conduit artery stiffness. *Journal of Biomedical Optics*, 18(2):027004, feb 2013. ISSN 1083-3668. doi: 10.1117/1.JBO.18.2.027004. URL http://biomedicaloptics.spiedigitallibrary.org/article.aspx?doi=10.1117/1.JBO.18.2.027004.

[15] C. Harris and M. Stephens. A Combined Corner and Edge Detector. In *Procedings of the Alvey Vision Conference 1988*, volume 69, pages 23.1–23.6. Alvey Vision Club, 1988. ISBN 604450133. doi: 10.5244/C.2.23. URL http://www.bmva.org/bmvc/1988/avc-88-023.html.

[16] Xiaochuan He, Rafik A. Goubran, and Xiaoping P. Liu. Wrist pulse measurement and analysis using Eulerian video magnification. *3rd IEEE EMBS International Conference on Biomedical and Health Informatics, BHI 2016*, pages 41–44, 2016. doi: 10.1109/BHI.2016.7455830.

[17] C. R. Hertzman, A. B. and Spealman. Observations on the finger volume pulse recorded photoelectrically. *American Journal of Physiology*, 119:334–335, 1937.

[18] A.A.R. Kamal, J.B. Harness, G. Irving, and A.J. Mearns. Skin photoplethysmography — a review. *Computer Methods and Programs in Biomedicine*, 28(4):257–269, apr 1989. ISSN 01692607. doi: 10.1016/0169-2607(89)90159-4. URL https://linkinghub.elsevier.com/retrieve/pii/0169260789901594.

[19] Syed Ghufran Khalid, Jufen Zhang, Fei Chen, and Dingchang Zheng. Blood Pressure Estimation Using Photoplethysmography Only: Comparison between Different Machine Learning Approaches. *Journal of Healthcare Engineering*, 2018, 2018. ISSN 20402309. doi: 10.1155/2018/1548647.

[20] Toshio Koga. Motion compensated interframe coding for video-conferencing. In *Proc. Nat. Telecommun. Conf.*, pages G5—-3, 1981.

[21] Xiaobai Li, Jie Chen, Guoying Zhao, and Matti Pietikäinen. Remote heart rate measurement from face videos under realistic situations. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 4264–4271, 2014. ISSN 10636919. doi: 10.1109/CVPR.2014.543.

[22] L. G. Lindberg and P Å. Öberg. Photoplethysmography. *Medical & Biological Engineering & Computing*, 29(1):48–54, jan 1991. ISSN 0140-0118. doi: 10.1007/BF02446295. URL http://ukpmc.ac.uk/abstract/MED/2016920http://link.springer.com/10.1007/BF02446295.

[23] Lars-Goeran Lindberg. Optical properties of blood in motion. *Optical Engineering*, 32(2):253, 1993. ISSN 00913286. doi: 10.1117/12.60688. URL http://opticalengineering.spiedigitallibrary.org/article.aspx?doi=10.1117/12.60688.

[24] Litong Feng, Lai-Man Po, Xuyuan Xu, Yuming Li, and Ruiyi Ma. Motion-Resistant Remote Imaging Photoplethysmography Based on the Optical Properties of Skin. *IEEE Transactions on Circuits and Systems for Video Technology*, 25(5):879–891, may 2015. ISSN 1051-8215. doi: 10.1109/TCSVT.2014.2364415. URL http://ieeexplore.ieee.org/document/6933875/.

[25] Le Liu, Le Lu, Jingjing Luo, Jun Zhang, and Xiuhong Chen. Enhanced Eulerian video magnification. *Proceedings - 2014 7th International Congress on Image and Signal Processing, CISP 2014*, pages 50–54, 2014. doi: 10.1109/CISP.2014.7003748.

[26] D.G. Lowe. Object recognition from local scale-invariant features. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, pages 1150–1157 vol.2. IEEE, 1999. ISBN 0-7695-0164-8. doi: 10.1109/ICCV.1999.790410. URL http://ieeexplore.ieee.org/document/790410/.

[27] Bruce D Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision, 1981.

[28] Christopher J. Marrocco and H. Ruth L. Bush. Peripheral arterial disease. *High Risk Diabetic Foot: Treatment and Prevention*, 358:1–8, 2010. doi: 10.3109/9781420083026.

[29] Daniel J. McDuff, Justin R. Estepp, Alyssa M. Piasecki, and Ethan B. Blackford. A survey of remote optical photoplethysmographic imaging methods. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, volume 2015-Novem, pages 6398–6404. IEEE, aug 2015. ISBN 978-1-4244-9271-8. doi: 10.1109/EMBC.2015.7319857. URL http://ieeexplore.ieee.org/document/7319857/.

[30] Jermana Moraes, Matheus Rocha, Glauber Vasconcelos, José Vasconcelos Filho, Victor de Albuquerque, and Auzuir Alexandria. Advances in Photopletysmography Signal Analysis for Biomedical Applications. *Sensors*, 18(6):1894, jun 2018. ISSN 1424-8220. doi: 10.3390/s18061894. URL http://www.mdpi.com/1424-8220/18/6/1894.

[31] V.S. Murthy, Sripad Ramamoorthy, Narayanan Srinivasan, Sriram Rajagopal, and M.M. Rao. Analysis of photoplethysmographic signals of cardiovascular patients. In *2001 Conference Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 3, pages 2204–2207. IEEE, 2001. ISBN 0-7803-7211-5. doi: 10.1109/IEMBS.2001.1017209. URL http://ieeexplore.ieee.org/document/1017209/.

[32] Salvatore Novo. Classification, epidemiology, risk factors, and natural history of peripheral arterial disease. *Diabetes, Obesity and Metabolism*, 4(s2):S1–S6, mar 2002. ISSN 1462-8902. doi: 10.1046/j.1463-1326.2002.0040s20s1.x. URL http://doi.wiley.com/10.1046/j.1463-1326.2002.0040s20s1.x.

[33] Tae Hyun Oh, Ronnachai Jaroensri, Changil Kim, Mohamed Elgharib, Frédo Durand, William T. Freeman, and Wojciech Matusik. Learning-Based Video Motion Magnification. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11208 LNCS:663–679, 2018. ISSN 16113349. doi: 10.1007/978-3-030-01225-0_39.

[34] Philipp V. Rouast, Marc T. P. Adam, Raymond Chiong, David Cornforth, and Ewa Lux. Remote heart rate measurement using low-cost RGB face video: a technical literature review. *Frontiers of Computer Science*, 12(5):858–872, oct 2018. ISSN 2095-2228. doi: 10.1007/s11704-016-6243-6. URL http://link.springer.com/10.1007/s11704-016-6243-6.

[35] U. Rubins, R. Erts, and V. Nikiforovs. The blood perfusion mapping in the human skin by photoplethysmography imaging. In *IFMBE Proceedings*, volume 29, pages 304–306. 2010. ISBN 9783642130380. doi: 10.1007/978-3-642-13039-7_76. URL http://link.springer.com/10.1007/978-3-642-13039-7{_}76.

[36] Abhijit Sarkar, A. Lynn Abbott, Zachary Doerzaph, and Kayla Sykes. Evaluation of video magnification for nonintrusive heart rate measurement. In *2016 IEEE First International Conference on Control, Measurement and Instrumentation (CMI)*, number Cmi, pages 494–498. IEEE, jan 2016. ISBN 978-1-4799-1769-3. doi: 10.1109/CMI.2016.7413797. URL http://ieeexplore.ieee.org/document/7413797/.

[37] Yu Sun and Nitish Thakor. Photoplethysmography Revisited: From Contact to Noncontact, From Point to Imaging. *IEEE Transactions on Biomedical Engineering*, 63(3):463–477, mar 2016. ISSN 0018-9294. doi: 10.1109/TBME.2015.2476337. URL http://ieeexplore.ieee.org/document/7268900/.

[38] Yu Sun, Charlotte Papin, Vicente Azorin-Peris, Roy Kalawsky, Stephen Greenwald, and Sijung Hu. Use of ambient light in remote photoplethysmographic systems: comparison between a high-performance camera and a low-cost webcam. *Journal of Biomedical Optics*, 17(3):037005, 2012. ISSN 10833668. doi: 10.1117/1.JBO.17.3.037005. URL http://biomedicaloptics.spiedigitallibrary.org/article.aspx?doi=10.1117/1.JBO.17.3.037005.

[39] Sungjun Kwon, Jeehoon Kim, Dongseok Lee, and Kwangsuk Park. ROI analysis for remote photoplethysmography on facial video. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, volume 2015-Novem, pages 4938–4941. IEEE, aug 2015. ISBN 978-1-4244-9271-8. doi: 10.1109/EMBC.2015.7319499. URL http://ieeexplore.ieee.org/document/7319499/.

[40] Wim Verkruysse, Lars O Svaasand, and J Stuart Nelson. Remote plethysmographic imaging using ambient light. *Optics Express*, 16(26):21434, dec 2008. ISSN 1094-4087. doi: 10.1364/OE.16.021434. URL https://www.osapublishing.org/abstract.cfm?URI=oe-16-26-21434.

[41] Karen Visser and M. G. Myriam Hunink. Peripheral Arterial Disease: Gadolinium-enhanced MR Angiography versus Color-guided Duplex US—A Meta-analysis. *Radiology*, 216(1):67–77, jul 2000. ISSN 0033-8419. doi: 10.1148/radiology.216.1.r00jl0367. URL http://pubs.rsna.org/doi/10.1148/radiology.216.1.r00jl0367.

[42] Neal Wadhwa, Michael Rubinstein, Frédo Durand, and William T. Freeman. Phase-based video motion processing. *ACM Transactions on Graphics*, 32(4):1, jul 2013. ISSN 07300301. doi: 10.1145/2461912.2461966. URL http://dl.acm.org/citation.cfm?doid=2461912.2461966.

[43] Neal Wadhwa, Michael Rubinstein, Fredo Durand, and William T. Freeman. Riesz pyramids for fast phase-based video magnification. In *2014 IEEE International Conference on Computational Photography (ICCP)*, pages 1–10. IEEE, may 2014. ISBN 978-1-4799-5188-8. doi: 10.1109/ICCPHOT.2014.6831820. URL http://ieeexplore.ieee.org/document/6831820/.

[44] Wenjin Wang. *Robust And Automatic Remote Photoplethysmography*. Number 2017. 2017. ISBN 978-90-386-4359-5. URL https://pure.tue.nl/ws/files/78340965/20171023{_}Wang.pdf{%}0Ahttps://research.tue.nl/en/publications/robust-and-automatic-remote-photoplethysmography{%}0Ahttps://pure.tue.nl/ws/portalfiles/portal/78340965/20171023{_}Wang.pdf.

[45] Wenjin Wang, Bert Den Brinker, Sander Stuijk, and Gerard De Haan. Algorithmic principles of remote-PPG. 64(7):1479–1491, 2019. doi: 10.1109/TBME.2016.2609282.

[46] Wenjin Wang, Albertus C. den Brinker, and Gerard de Haan. Single-Element Remote-PPG. *IEEE Transactions on Biomedical Engineering*, 66(7):2032–2043, jul 2019. ISSN 0018-9294. doi: 10.1109/TBME.2018.2882396. URL https://ieeexplore.ieee.org/document/8540913/.

[47] Hao-Yu Wu, Michael Rubinstein, Eugene Shih, John Guttag, Frédo Durand, and William Freeman. Eulerian video magnification for revealing subtle changes in the world. *ACM Transactions on Graphics*, 31(4):1–8, jul 2012. ISSN 07300301. doi: 10.1145/2185520.2335416. URL http://dl.acm.org/citation.cfm?doid=2185520.2335416.

[48] Xiu Wu, Xuezhi Yang, Jing Jin, and Zhao Yang. PCA-based magnification method for revealing small signals in video. *Signal, Image and Video Processing*, 12(7):1293–1299, 2018. ISSN 18631711. doi: 10.1007/s11760-018-1282-0. URL https://doi.org/10.1007/s11760-018-1282-0.

[49] Matthew F. Wyatt, Chad Stickrath, Alap Shah, Alexandra Smart, James Hunt, and Ivan P. Casserly. Ankle—brachial index performance among internal medicine residents. *Vascular Medicine*, 15 (2):99–105, apr 2010. ISSN 1358-863X. doi: 10.1177/1358863X09356015. URL http://journals.sagepub.com/doi/10.1177/1358863X09356015.

[50] K Qawqzeh Yousef, Uldis Rubins, and Alharbi Mafawez. Photoplethysmogram second derivative review: Analysis and applications. *Scientific Research and Essays*, 10(21):633–639, nov 2015. ISSN 1992-2248. doi: 10.5897/SRE2015.6322. URL http://academicjournals.org/journal/SRE/article-abstract/A214BA956534.

[51] Sebastian Zaunseder, Alexander Trumpp, Daniel Wedekind, and Hagen Malberg. Cardiovascular assessment by imaging photoplethysmography – a review. *Biomedical Engineering / Biomedizinische Technik*, 63(5):617–634, oct 2018. ISSN 1862-278X. doi: 10.1515/bmt-2017-0119. URL http://www.degruyter.com/view/j/bmte.2018.63.issue-5/bmt-2017-0119/bmt-2017-0119.xml.

[52] Yichao Zhang, Silvia L. Pintea, and Jan C. van Gemert. Video Acceleration Magnification. In *Computer Vision and Pattern Recognition*, 2017. URL https://acceleration-magnification.github.io/.