Detecting stress from imaging photoplethysmography using high frame rate video and a yellow-green filter: A pilot study

P. Aquilina[1] D. Booth[2] G. Hanly[3] K. Meaney[4]
S. Darvishi[5] B. M. Pincombe[6]

(Received 29 February 2020; revised 14 September 2020)

Abstract

We investigate the use of a yellow-green filter to increase the signal-to-noise ratio (SNR) in imaging photoplethysmography (iPPG) and test if high frame rate (HFR) video improves the accuracy of the derived heart rate variability (HRV). This pilot study is associated with a broader program to use iPPG to detect and monitor stress levels using HRV. To improve the SNR of the iPPG signal, we employ two HFR colour video cameras of which one was fitted with a yellow-green filter (corresponding to the haemoglobin absorption peak within the visible spectrum). To our knowledge, the benefit of a yellow-green filter has never been explored. The predominant influence on HRV comes from

DOI:10.21914/anziamj.v61i0.15186, Published 2020-09-26, as part of the Proceedings of the 14th Biennial Engineering Mathematics and Applications Conference. ISSN 1445-8810. (Print two pages per sheet of paper.)

the autonomic nervous system (ANS), which connects directly to the heart and cues the human body to relax or to stress. The linkage of HRV to the ANS makes HRV a proxy for stress levels. The HRV is derived from the iPPG signal by first using a cubic spline interpolation for more precise peak detection, and then calculating the inter-beat intervals from the peak-to-peak time differences. Instead of interpolating the signal, we hypothesise that a more accurate HRV measurement can be obtained using a HFR video camera, in our case at 200 frames per second.

Contents

1 Introduction **C275**
1.1 Motivation C275

2 Background **C276**
2.1 Imaging photoplethysmography (iPPG) C276
2.2 Heart rate variability (HRV) C277

3 Data collection **C278**
3.1 Why high frame rate video? C279
3.2 Why a yellow-green filter? C280
3.3 The setup C281

4 Results and analysis **C281**
4.1 High frame rate video C282
4.2 Yellow-green filter C283

5 Future work **C284**

6 Conclusion **C285**

1 Introduction

This pilot study examines the feasibility of detecting stress in an individual based on their heart rate variability (HRV). Individual heart beats are detected by applying imaging photoplethysmography (iPPG) techniques (also known as remote photoplethysmography or photoplethysmography imaging) to video collected by a pair of high frame rate (HFR, 200 Hz) cameras. The particular emphasis of this study is in assessing the benefits of using a HFR camera and a particular yellow-green (575 nm) filter which was chosen because it corresponds to the absorption peak of haemoglobin. The HRV is one of the three classical markers of autonomic nervous system (ANS) activation, alongside galvanic skin response and pupil dilation. Traditionally HRV has been gathered using contact measures but iPPG is preferable because of the motivating problem of stress detection in officers. This pilot study identifies means to improve the iPPG signal quality for use in a subsequent larger study comparing multiple measures of stress on many human participants.

1.1 Motivation

This project grew from a conversation that a colleague had with Australian Federal Police (AFP) officers who view child exploitation imagery for investigative purposes. Viewing exploitation imagery is often distressing to and has consequences for the viewer's physical health, relationships, trust and mental health. The AFP already has processes in place to protect these officers. They include rotation to an alternative position after a predetermined length of time, regardless of the individual's assessment of their own wellbeing, and regardless of whether or not they have adopted viewing protocols to reduce the potency of the imagery. These protocols are similar to those given to journalists to reduce harm from disturbing images. They are mostly self-imposed suggestions such as turn audio down or ideally off, take regular breaks and exercise [2].

Therefore, our true motivation, is to develop an external automated method for measuring stress. We would like this to be entirely non-contact and ideally

use readily available, inexpensive hardware such as a standard webcam.

2 Background

2.1 Imaging photoplethysmography (iPPG)

The technology with which we hope to realise our goal is known as photoplethysmography or PPG for short. This is where we measure the blood flow under the skin from either reflected or transmitted light. The signal that we get back has a DC or non-varying component due to the light source, skin and blood in the veins. But the part we are much more interested in is the varying or AC component of the signal. This provides a measure of the flow of blood in the arteries of the hypodermis—which is a layer connecting the skin to the bone and muscles.

A Fitbit (or any similar activity tracker/smartwatch) shines typically green light onto a wrist and measures whatever light comes back. This is contact PPG because the Fitbit is in contact with the body. For our application, we do not want to further burden the officers with wires or wearable sensors. We want to sense or capture this light from a distance. That is what video cameras do: they measure light from a distance. Using a video camera to extract the PPG signal is commonly referred to as imaging photoplethysmography. It is very convenient, but the signal is also noisy and challenging to work with. The noise is to be expected because we are attempting to extract the colour change from a patch of skin from video when we cannot even notice the change with the naked eye.

Here we explain how we convert video of a face to a waveform showing the blood volume under the skin (iPPG). First we need a way of choosing a patch of skin. The simplest method, is to manually crop a section of the skin on the first frame, and then just look in that same spot across all frames. But if the subject is moving, then more advanced image processing with tracking or face detection is called for. Regardless of which method is employed, skin

detection is usually helpful.

So, for every frame in the video we pull out our region of interest (ROI), which is just a collection of skin pixels. All these skin pixels are telling us the same thing—that is, how much blood is under the skin at that moment in time, so we average them to get a consensus. We do this averaging separately for each channel (we are using an RGB video camera), so that per frame we have a single number for red, a number for green and a number for blue.

We then apply some pre-processing. In our case, we detrend the signal using mean-centering-and-scaling.

Next we need to combine the colour signals in a way to extract the best possible PPG signal. We have chosen to use independent component analysis (ICA), which acknowledges that the multiple inputs are a tangled mess due to various sources. Some sources of noise include the blood volume signal, the movement noise, environmental noise and camera noise. With ICA the inputs are detangled by separating them into independent non-Gaussian signals. We then choose the source signal that looks most like blood volume by comparing the frequencies to the expected range of a human heart beat. Finally, we do some signal post-processing to clean it up; specifically, we apply a moving average filter and remove outliers.

2.2 Heart rate variability (HRV)

Figure 1 shows a very clean PPG signal with four heart beats. The time between each beat—the interbeat interval (IBI)—is different. This is normal. Our heart rate is not constant—people typically have considerable HRV, even between consecutive beats.

The HRV is largely caused by the autonomic nervous system (ANS). The ANS provides a direct connection between our brains and our hearts (and other organs) through a network of nerves. The job of the ANS is to regulate various bodily functions, typically without any conscious control or effort—hence the term autonomic. The ANS is composed of three subsystems but the two that

Figure 1: A typical example of how the length of time between each heart beat varies significantly.

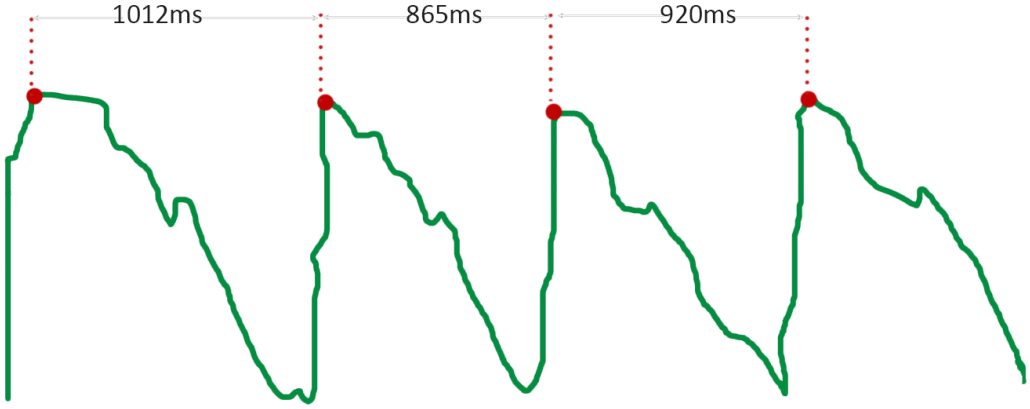

have the most control over HRV are the parasympathetic nervous system (PNS) and the sympathetic nervous system (SNS). The sympathetic nervous system is associated with stress responses. Under stressful situations the SNS activity increases so that our bodies are ready for a "flight-or-fight" response.

The PNS usually takes an opposing role to the SNS and is dominant when people are relaxing or digesting. The SNS can effect changes to one's heart rate much quicker than the PNS. Therefore, frequency analysis of HRV can provide insights into the relative activity of the PNS versus the SNS, and in turn gives insight into one's level of stress.

3 Data collection

What makes this work novel is the data we are collecting and analysing. We are investigating improvements, if any, in IBI measurements due to HFR video, and in the signal to noise ratio (SNR) of the iPPG due to a camera fitted with a narrow band yellow-green filter.

3.1 Why high frame rate video?

Each frame of the video provides one data point in the PPG signal. In our case with a 200 Hz video we have nearly seven times more points than a standard 30 Hz camera.

The high frequency is important because the HRV is measured in milliseconds and we need to precisely pinpoint the peak of each heart beat. Interpolating between points with a cubic spline (which is common practice in the literature [e.g., 4, 5]) helps to estimate the precise location of the peaks but cannot be as accurate as actually sampling at a higher density. One of our aims is to quantify the benefits of a higher frame rate.

We are not the first to consider if HFR video improves HRV estimation. In 2013, Sun et al. [4] compared the HRV derived from video at 200 Hz with the same video downsampled to 100 Hz, 50 Hz and 20 Hz (all subsequently interpolated back up to 200 Hz using the cubic spline). They concluded that the results were statistically similar regardless of the frame rate. We feel that it is worth revisiting HFR video for two reasons.

First, the study by Sun et al. [4] was a relatively small preliminary study (ten participants, each giving one sample of four minutes). Despite the work being published in 2013, we are not aware of any other work (by that group or any researcher) which verifies their preliminary findings relating to HFR video.

Secondly, Sun et al. [4] utilised a highly controlled setup that approached contact PPG. The participants sat perfectly still and further restrained movement in their left hand by resting it on a pillow. This allowed their palm to be recorded by a monochrome camera positioned just 400 mm away. The palm was also actively lit by two commercial infrared light sources. It is uncertain how well their conclusions would map to a more typical iPPG collection using a RGB video camera at least 1.5 m from the subject (as we have done).

Figure 2: Some wavelengths are better than others at recovering the PPG signal [1].

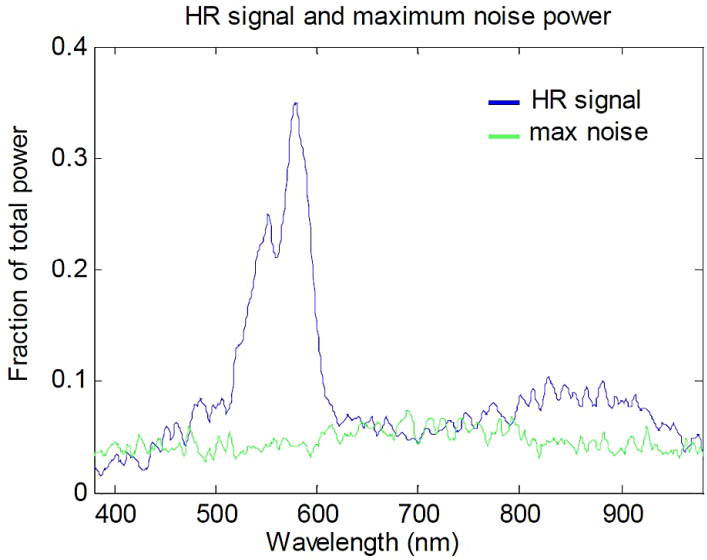

3.2 Why a yellow-green filter?

Some wavelengths are better than others at recovering the PPG signal, which is why Fitbit and other similar devices usually use green light. In an experiment looking at backscattered signals across a wide range of wavelengths, Martinez et al. [3] found that wavelengths with the best SNR corresponded to the absorption spectrum of haemoglobin. These results were confirmed by Blackford et al. [1] in a carefully controlled study.

The best wavelengths for PPG are in the band from 480 to 610 nm with the peak around 575 nm (see Figure 2). So one of the aims of this pilot study is to test if a yellow-green filter centered at 575 nm with a bandwidth of 27 nm will enable PPG to be extracted with a higher SNR.

3.3 The setup

For our experimental setup (see Figure 3) we chose two professional Sony video cameras (Sony PXW-SF7 MK2). They provide full HD video at 200 frames per second in a RAW format. One of the cameras was fitted with a yellow-green filter. The filter reduces the light throughput by a quarter and this was taken into account when selecting the camera. The cameras are placed as close as practicable to each other on a single tripod mount.

We also have four video studio lights with dynamic intensity and colour temperature control to provide good, controlled lighting.

The subject wears an electroencephalography (EEG) cap and there is a PPG ear clip to gather ground truth. These sensors reduce the exposed skin which can make iPPG recovery more challenging. The EEG and contact PPG data were not utilised in this pilot study. The intent is for this data to be analysed in a subsequent study looking at the efficacy of stress detection from EEG and PPG.

The cameras are both initiated simultaneously via remote control. The synchronisation error did not exceed 50 ms. The participant is seated before stimulus videos intended to either invoke no reaction or a stress response. A second monitor is placed within the video frame for contextual and synchronisation purposes.

4 Results and analysis

As Unakafov [5] noted, we also found that it was more precise to derive the IBIs from the minima between heartbeats rather than the maxima. In the results that follow, both maxima and minima are shown for illustration.

Figure 3: The experimental setup.

4.1 High frame rate video

As already mentioned, a high frame rate camera was utilised to test for any gains in accuracy of the IBI measurements.

The plot in Figure 4 demonstrates qualitatively the benefit of HFR video. The red line represents iPPG captured at approximately 30 Hz (standard for a webcam), emulated by downsampling the HFR video. In contrast, the blue line is the HFR-derived iPPG interpolated up to 400 Hz using the cubic spline. The right-most minima clearly shows that, due to the sampling sparsity, the 30 Hz misses the bottom of the pulse and so results in an inaccurate measurement.

Figure 4: The right-most minima shows that the 30 Hz signal misses the bottom of the pulse and so results in an inaccurate measurement.

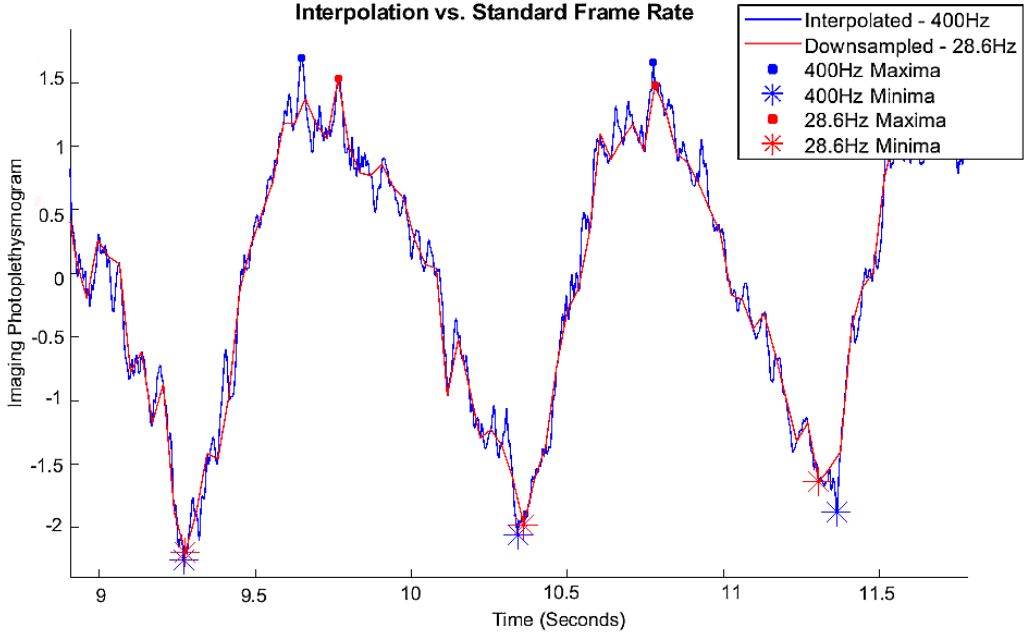

4.2 Yellow-green filter

Figure 5 shows decisively that the yellow-green filter has more noise than the iPPG derived from the full colour video. This is unfortunately the opposite result to our hypothesis and therefore the filter does not seem to assist in deriving a more accurate iPPG signal.

We believe this result is due to the limited light that was able to reach the sensor. There are losses of light as the light passes through the glass of the filter, and also due to the limited filter size. We were only able to procure a 50 mm filter off-the-shelf, which meant we had to custom make an adapter to fit to our 90 mm lens. These two factors together significantly reduced the number of photons that could reach the sensor.

Figure 5: Clearly the signal from the yellow-green filter (bottom) has more noise than the iPPG derived from the full colour video (top).

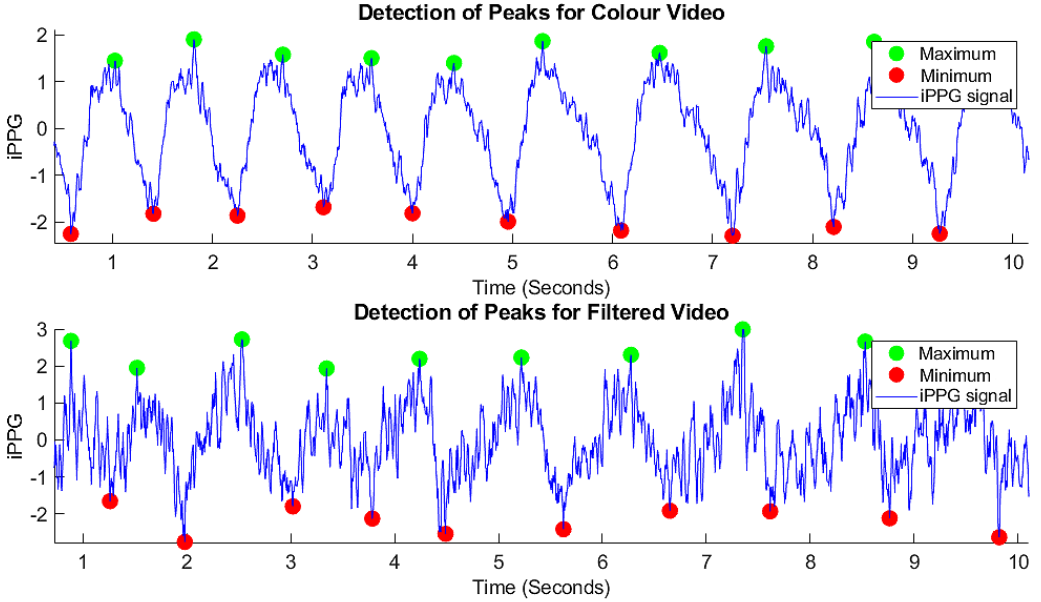

5 Future work

This article lays the groundwork for a larger data collection and a more rigourous investigation into the impacts of high frame rate video and the yellow-green filter.

For the purpose of drawing conclusions about the predictability of the methods described within this article for measuring stress, a much larger data collection is required.

Further investigation is needed to determine the merits of using a yellow-green filter. The setup will need to be altered to increase the light hitting the sensor and the benefit of fusing the yellow-green data with RGB data should be explored.

6 Conclusion

The preliminary results that are documented in this article suggest that the use of a yellow-green filter is not delivering on our expectations. It has not improved the quality of the iPPG signal. In fact, the iPPG derived from the yellow-green filtered video exhibited significantly greater noise than that derived from the non-filtered camera, likely as a consequence of the reduced amount of light reaching the sensor.

However, we have also demonstrated that high frame rate video can improve the accuracy of peak detection within iPPG. Given that HRV is only concerned with the relative changes in adjacent IBIs, small improvements in IBI measurements can be significant. The authors stress that the findings in this article are preliminary only and are to be validated in a larger, follow-up study.

Acknowledgements Thanks to our partners in the broader stress project Rebecca Heyer, Carolyn Semmler, Reg Nixon, Melanie Takarangi, Siobhan Banks, Sau Yee Yiu, Desmond Yau and Neil Gordon and to The Defence Innovation Partnership.

References

[1] E. B. Blackford, J. R. Estepp, and D. J. McDuff. "Remote spectral measurements of the blood volume pulse with applications for imaging photoplethysmography". In: *Optical Diagnostics and Sensing XVIII: Toward Point-of-Care Diagnostics*. Ed. by G. L. Coté. Vol. 10501. International Society for Optics and Photonics. SPIE, 2018, 105010Z. DOI: 10.1117/12.2291073 (cit. on p. C280).

[2] M. Brayne. *Trauma and Journalism: A Guide For Journalists, Editors and Managers*. DART Center for Journalism and Trauma, 2007. URL: https://dartcenter.org/sites/default/files/DCE_JournoTraumaHandbook.pdf (cit. on p. C275).

[3] L. F. Corral Martinez, G. Paez, and M. Strojnik. "Optimal wavelength selection for noncontact reflection photoplethysmography". In: *Proceedings of the 22nd Congress of the International Commission for Optics: Light for the Development of the World*. Vol. 8011. International Society for Optics and Photonics. SPIE, 2011, p. 801191. DOI: 10.1117/12.903190 (cit. on p. C280).

[4] Y. Sun, S. Hu, V. Azorin-Peris, R. Kalawsky, and S. E. Greenwald. "Noncontact imaging photoplethysmography to effectively access pulse rate variability". In: *J. Biomed. Optics* 18.6 (2013), p. 061205. DOI: 10.1117/1.JBO.18.6.061205 (cit. on p. C279).

[5] A. M. Unakafov. "Pulse rate estimation using imaging photoplethysmography: generic framework and comparison of methods on a publicly available dataset". In: *Biomed. Phys. Eng. Exp.* 4.4 (2018), p. 045001. DOI: 10.1088/2057-1976/aabd09 (cit. on pp. C279, C281).

Author addresses

1. **P. Aquilina**, 205 Labs, Biometrics Group, IA, ISSD, DST Group, Edinburgh, South Australia 5111, AUSTRALIA.
mailto:peter.aquilina@dst.defence.gov.au
orcid:https://orcid.org/0000-0001-5399-1299

2. **D. Booth**, 205 Labs, Biometrics Group, IA, ISSD, DST Group, Edinburgh, South Australia 5111, AUSTRALIA.
mailto:david.booth@dst.defence.gov.au

3. **G. Hanly**, 205 Labs, Biometrics Group, IA, ISSD, DST Group, Edinburgh, South Australia 5111, AUSTRALIA.
mailto:gary.hanly@dst.defence.gov.au

4. **K. Meaney**, 205 Labs, Biometrics Group, IA, ISSD, DST Group, Edinburgh, South Australia 5111, AUSTRALIA.
mailto:Kym.Meaney2@dst.defence.gov.au

5. **S. Darvishi**, ElectroAutoMedics Pty Ltd, 296 Glynburn Rd, Kensington Gardens, South Australia 5068, Australia.
mailto:sam.darvishi@electroautomedics.com

6. **B. M. Pincombe**, 2.F.46, 205 Labs, Social and Behavioural Sciences Group, JWO, JOAD, DST Group, Edinburgh, South Australia 5111, Australia.
mailto:brandon.pincombe@dst.defence.gov.au