



Universidad Autónoma  
de Madrid

**Biblos-e Archivo**  
Repositorio Institucional UAM

**Repositorio Institucional de la Universidad Autónoma de Madrid**  
<https://repositorio.uam.es>

Esta es la **versión de autor** del artículo publicado en:  
This is an **author produced version** of a paper published in:

J. Hernandez-Ortega, J. Fierrez, A. Morales and D. Diaz, "A Comparative Evaluation of Heart Rate Estimation Methods using Face Videos," *2020 IEEE 44th Annual Computers, Software, and Applications Conference (COMPSAC)*, 2020, pp. 1438-1443

**DOI:** <http://doi.org/10.1109/COMPSAC48688.2020.00-53>

**Copyright:** © 2020 Institute of Electrical and Electronics Engineers

El acceso a la versión del editor puede requerir la suscripción del recurso  
Access to the published version may require subscription

# A Comparative Evaluation of Heart Rate Estimation Methods using Face Videos

Javier Hernandez-Ortega, Julian Fierrez, Aythami Morales, David Diaz  
School of Engineering, Universidad Autonoma de Madrid, Spain  
{javier.hernandezo, julian.fierrez, aythami.morales}@uam.es

**Abstract**—This paper presents a comparative evaluation of methods for remote heart rate estimation using face videos, i.e., given a video sequence of the face as input, methods to process it to obtain a robust estimation of the subject’s heart rate at each moment. Four alternatives from the literature are tested, three based in hand-crafted approaches and one based on deep learning. The methods are compared using RGB videos from the COHFACE database. Experiments show that the learning-based method achieves much better accuracy than the hand-crafted ones. The low error rate achieved by the learning-based model makes possible its application in real scenarios, e.g. in medical or sports environments.

**Index Terms**—Remote Plethysmography, Heart Rate, Face Biometrics, Hand-crafted, Deep Learning

## I. INTRODUCTION

Photoplethysmography (PPG) [1] is a low-cost technique for measuring the cardiovascular Blood Volume Pulse (BVP) through changes in the amount of light reflected or absorbed by human vessels. This information can be used to estimate parameters such as heart rate, arterial pressure, blood glucose level, or oxygen saturation levels. PPG signals are usually measured with contact sensors placed at the fingertips, the chest, or the feet. This type of contact measurement may be suitable for a clinic environment, but it can be uncomfortable and inconvenient for other application fields such as driver monitoring, sport events, or face antispoofing. Remote PPG (rPPG) consists in applying PPG techniques to face video sequences. These techniques look for changes in the color of the user’s face caused by the variation of the oxygen concentration in the blood.

Classic rPPG works, usually called hand-crafted rPPG, use signal processing techniques for analyzing the images, and looking for slight color and illumination changes related with the BVP. More recent approaches use learning-based alternatives in order to train models capable of estimating the heart rate when receiving a face video in their input. A selection of works related to remote photoplethysmography is shown in Table I.

One of the drawbacks of heart rate estimation using rPPG, compared to contact-based PPG techniques, is that rPPG is very sensitive to variability in the acquisition scenario. Some sources of this variability are [14]: changes in the illumination, blur, movement of the subject, vibration of the camera, resolution, frame rate, etc. Under constrained conditions, hand-crafted methods are able to achieve high accuracy levels, but



Fig. 1. Example images from the COHFACE database for its two illumination scenarios: (a) controlled illumination, and (b) uncontrolled illumination.

in unconstrained scenarios learning-based approaches appear to dominate.

The target of the present work is to perform a comparative study of state-of-the-art rPPG methods for heart rate estimation. We compare different methods presented in the literature, both hand-crafted and learning-based, tested against the same selection of videos from the COHFACE database [15]. We check the performance of both types of approaches, identifying their respective strengths and weaknesses.

The rest of the paper is structured as follows. In Sect. II we describe the database, the methods, and the experimental protocol. The results of the experiments are reported in Sect. III. Conclusions and future work are finally drawn in Sect. IV.

## II. EXPERIMENTAL FRAMEWORK

### A. Database

The database used for the experiments is COHFACE [15]. This database is composed of 160 videos accompanied by the corresponding groundtruth heart rates of the subjects. In total 40 people participated and 4 videos were obtained from each of them. The duration of the recorded videos is 1 minute each. The physiological signals relative to the heart rate were obtained by contact photoplethysmography (PPG). Both signals (video and PPG) are synchronized. The video sequences have a resolution of  $640 \times 480$  pixels and were captured with a sampling frequency of 20 frames per second. The heart rate signal was acquired at 256 Hz. Videos were recorded in a laboratory. Regarding illumination conditions, for each subject two videos were obtained with controlled lighting, and other two with natural lighting. Examples of the database are shown in Fig. 1.

TABLE I

SELECTION OF WORKS THAT USE DIFFERENT TYPES OF IMAGES TO IMPLEMENT rPPG FOR HEART RATE EXTRACTION OR OTHER RELATED TASKS LIKE FACE ANTISPOOFING OR STRESS DETECTION. THE WORKS EVALUATED IN THIS PAPER ARE HIGHLIGHTED IN BOLD.

Method	Type of Images	Database used	Video Length	Target	Performance
Garbey et al. 2007 [2]	Thermal	self-collected	120 secs.	HR Estimation	Accuracy = 99%
<b>Verkruyse et al. 2008 [3]</b>	RGB	self-collected	10 secs.	HR & Respiration Estimation	Qualitative
<b>Poh et al. 2011 [4]</b>	RGB	self-collected	60 secs.	HR Estimation	RMSE = 5.63%
<b>De Haan et al. 2013 [5]</b>	RGB	self-collected	2-30 secs.	HR Estimation	RMSE = 0.5 bpm
Tasli et al. 2014 [6]	RGB	self-collected	90 secs.	HR Estimation	MAE = 4.2 %
Chen et al. 2014 [7]	Hyperspectral	self-collected	30-60 secs.	Stress Estimation	Qualitative
McDuff et al. 2014 [8]	Multiband (RGBCO)	self-collected	120 secs.	HR Estimation	Correlation = 1.0
Chen et al. 2016 [9]	RGB + NIR	self-collected	90 secs.	HR Estimation	RMSE = 1.65%
Li et al. 2016 [10]	RGB	3DMAD and self-collected	10 secs.	Face Antispoofing	EER = 4.71%
<b>Wang et al. 2017 [11]</b>	RGB	self-collected	2-30 secs.	HR Estimation	SNR = 5.16
Hernandez et al. 2018 [12]	RGB and NIR	3DMAD and self-collected	1-60 secs.	HR Estimation	EER = 22% (RGB) — 0% (NIR)
<b>Chen et al. 2018 [13]</b>	RGB and NIR	HCI and self-collected	10 secs.	Face Antispoofing	MAE = 4.57 bpm

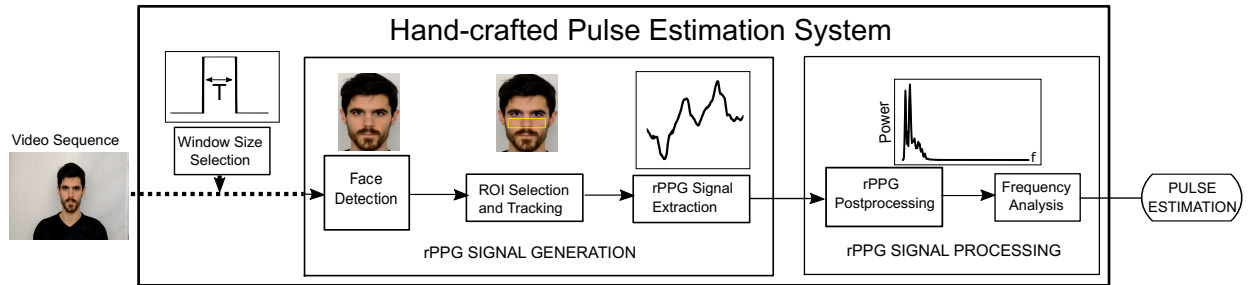


Fig. 2. Example of the workflow of a hand-crafted rPPG system for heart rate estimation using face videos. First, the different channels of the videos are processed to obtain raw rPPG signals. After that, the signals are postprocessed in order to extract a value for the heart rate for each temporal window.

We decided to use this database for the evaluation because of several reasons. First, it is one of the few publicly available datasets oriented to heart rate extraction. It also presents enough amount of data for being able to train a learning-based model from scratch. Finally, it also presents a good balance between favorable and challenging conditions with respect to frame rate, resolution and illumination.

### B. Methods under Evaluation

The methods compared in this paper can be split into two different categories: hand-crafted and learning-based.

**Hand-crafted methods** apply signal processing techniques to the video frames acquired by the camera sensors. These techniques analyze the video sequences and extract predefined features from the frames in order to estimate the heart rate signal. Hand-crafted methods have a major limitation: they may fail to model properly the highly nonlinear processes that may be occurring in the phenomenon they are trying to measure, as it happens in most of the processes related to human physiology, e.g. heart rate and respiration. An example of the typical workflow of a hand-crafted heart rate estimation method can be seen in Fig. 2.

The 4 *hand-crafted methods* that we compare in this work are: [3], which emphasizes the role of luminance for obtaining physiological information; [4], which uses Blind Source Separation through ICA; [5], which focuses on chrominance as the source of information; and [11], which uses what the authors

call The Plane Orthogonal to Skin—Tone (a linear function applied to the input data). Public open source implementations have been used for these 4 hand-crafted methods [16].

As a representative example that illustrates typical hand-crafted methods, we now sketch the method [4]. In this case the authors try to estimate the heart rate from a RGB video using Blind Source Separation (BSS) based on Independent Component Analysis (ICA), see Fig. 2:

- A facial Region Of Interest (ROI) is obtained for each frame using automatic face tracking. The ROI is separated into its three RGB channels and the pixels of each channel are spatially averaged to obtain their mean values in each frame. This way the raw rPPG signals  $x_1(t)$ ,  $x_2(t)$  and  $x_3(t)$  are generated by concatenating the mean values of all the processed frames. Subsequent processing is carried out using temporal windows of 30 seconds with a window overlap of the 96.7% (increments of 1 second). Normalization is applied to the raw rPPG signals in order to obtain signals with zero mean and unitary standard deviation.
- ICA is used to break down the normalized signals into three independent sources. Although the components obtained from ICA are not ordered, typically the second component contains a strong plethysmographic signal. Finally, the Fast Fourier Transform (FFT) is applied to the selected component in order to obtain its representation in the frequency domain. The final estimation of the heart

TABLE II

RESULTS OF THE COMPARISON OF HEART RATE ESTIMATION METHODS FROM FACE VIDEOS. MAE IS COMPUTED FOR 3 DIFFERENT VIDEO LENGTHS: 60, 30, AND 15 SECONDS. HIGHLIGHTED IN BOLD ARE THE BEST RESULTS OBTAINED HERE FOR EACH VIDEO DURATION. IN ITALIC ARE THE RESULTS PUBLISHED BY CHEN ET AL. 2018 [13].

MAE (bpm)	60 s	30 s	15 s	Mean
Verkruijse et al. 2008 [3]	12.86	14.30	16.91	14.69
Poh et al. 2011 [4]	16.51	13.75	14.90	15.05
De Haan et al. 2013 [5]	12.76	12.61	14.51	13.29
Wang et al. 2017 [11]	24.40	23.38	24.81	24.20
<b>Chen et al. 2018 [13] (COHFACE)</b>	<b>4.20</b>	<b>3.79</b>	<b>5.71</b>	<b>4.57</b>
<i>Chen et al. 2018 [13] (HCI)</i>	-	4.57	-	4.57

rate will be the frequency component that corresponds to the maximum peak of the spectrum inside the band of frequencies considered physiologically possible for human heart rate.

On the other hand, **learning-based methods** for heart rate estimation are of recent development. In our experimental comparison we include the *DeepPhys model* [13]. The main limitation of hand-crafted methods (i.e. its limited ability to model complex nonlinear functions) motivated the authors in [13] for designing a deep learning model that uses a Convolutional Attention Network (CAN) to perform physiological measurements from videos. This model achieves better results than traditional techniques since it is able of adapting its parameters to the non-linear nature of the heart rate signals. In addition, DeepPhys also permits the spatio-temporal visualization of physiological signals from RGB videos, e.g. the blood perfusion under the face skin.

### C. Experimental Protocol

We have split the COHFACE database into train, development, and test subsets. Of the 40 subjects in the database, 30 have been used for training, 1 has been used for validation and the remaining 9 have been used for testing. We had 120 minutes of video for training, 4 minutes for validation, and 36 minutes for testing.

First, for the hand-crafted methods, as they don't need training, the performance and error measures are directly obtained from the test subset of COHFACE. The learning-based model, i.e. DeepPhys, needs to be trained to learn how to estimate the heart rate from video sequences. The original model pretrained in [13] is not public, so we have replicated its architecture and we have trained it from scratch with the training and validation partitions of COHFACE.

Tests have been performed for video sequences of 15, 30, and 60 seconds of duration, and a sliding step of 1 second, in order to evaluate the possible application of the methods when different amounts of data are available. The estimated and the groundtruth heart rates have been compared in beats per minute (bpm). The metric chosen to check the performance of the models is the Mean Absolute Error (MAE), which evaluates the distance between the estimated heart rate obtained and the groundtruth measure.

## III. RESULTS

### A. Average Performance Comparison

Table II shows the results obtained with the hand-crafted and the learning-based approaches. It can be seen that for each method there are slight differences depending on the duration of the test videos. The mean MAE across the three lengths of the test videos is also presented.

Among the hand-crafted methods, the first three are performing similarly with mean MAE values around 15 bpm. However, the last one is obtaining worse error rates (around 24 bpm mean MAE). In that paper, the authors presented an alternative rPPG method where they defined a projection Plane Orthogonal to the Skin–Tone (POS). As they wanted to keep the algorithm as simple as possible, they did not use the common band-pass filtering usually employed in other rPPG methods. Due to its characteristics, the use of the POS algorithm is interesting in those cases with favorable conditions, but it is more vulnerable to external factors like a heterogeneous illumination spectra, i.e. when two or more illumination sources are present. The COHFACE videos used for testing contain both controlled and uncontrolled illumination conditions, and this is likely the reason of the poor performance of the method from [11] compared to the other hand-crafted alternatives. In general, the level of accuracy shown by the hand-crafted methods does not allow to apply them in many scenarios in which an error around 15 bpm may be too high.

The results in Table II also show that the DeepPhys model performs significantly better than the hand-crafted alternatives. This large difference in the performance is mainly caused by the ability of DeepPhys to apply the knowledge acquired from the training data, adapting its parameters to the conditions of each particular scenario. The learning-based model is capable of improving the hand-crafted results by around a 66% (from 15 bpm to 5 bpm of mean MAE). This is caused by its ability to capture the non-linear nature of the physiological data. In addition, the hand-crafted models have been designed to operate under certain conditions of illumination, resolution, etc, while the learning process allows to adapt the DeepPhys model to each specific scenario. On the other hand, the drawback of the learning-based approaches is that a high amount of data is needed in order to train a deep model from scratch, while hand-crafted methods do not need that training stage, being able to work directly on the test data. However, it is possible to reduce the necessary amount of training data using a model that has been previously trained with a similar dataset. That model can be then fine-tuned with few additional data to get adapted to the new scenario.

Regarding the video duration, increasing the length above 30 seconds does not improve the accuracy of the methods significantly. With video segments of at least 15 seconds the rPPG algorithms achieve their best results, so in case of wanting to apply this technology to a real scenario it seems that it will not be necessary to process longer videos.

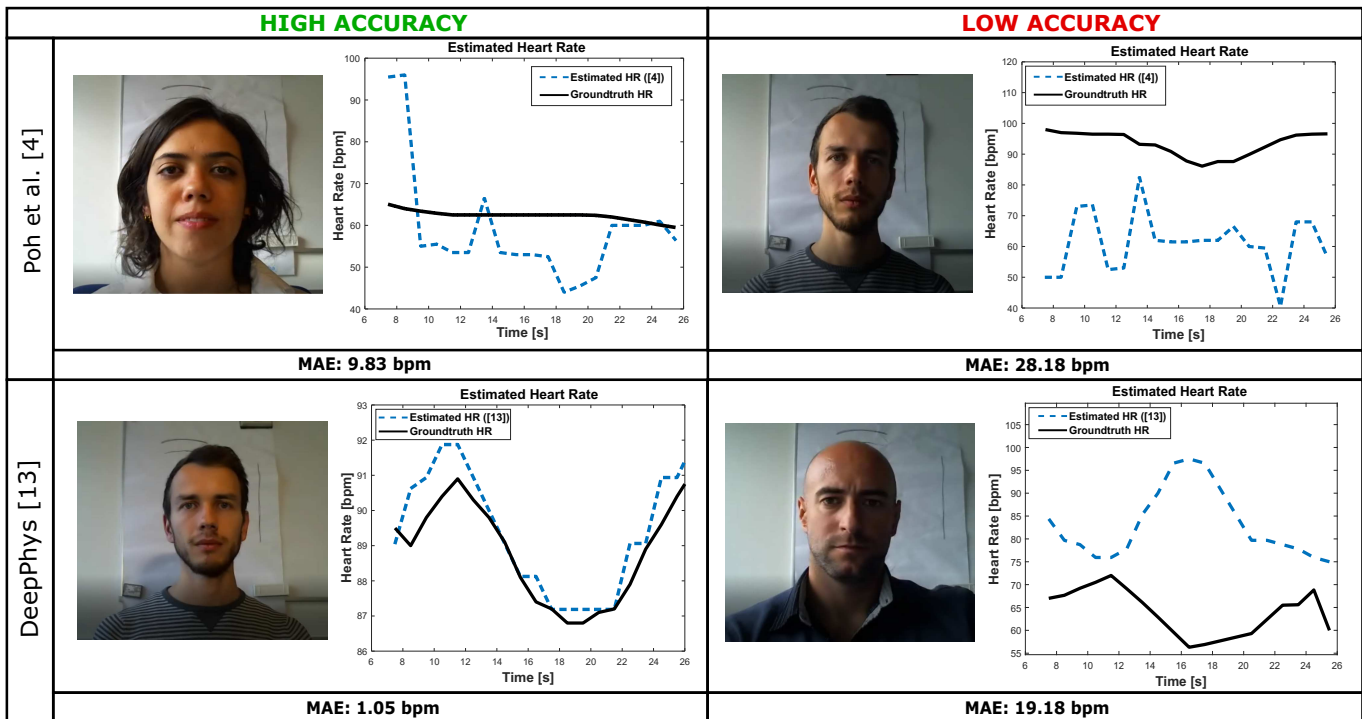


Fig. 3. **Accuracy for heart rate estimation** using a hand-crafted algorithm (Poh et al. [4]) and our model of DeepPhys trained with COHFACE. MAE is computed for a video length of 15 seconds, with a sliding step of 1 second. We have selected the videos that obtain the highest and the lowest accuracy for both methods. The performance is higher (lower MAE) when dealing with videos with controlled illumination and low movement (left column). Results are worse when the illumination is natural and unconstrained (right column).

We should also mention that in [13] the authors trained their DeepPhys model with videos from the Mahnob-HCI database [18] with a duration of 30 seconds, achieving a MAE value of 4.57 bpm, close to the performance obtained on average with the data we used in our experiments. It seems that the Convolutional Attention Network of DeepPhys is able to perform robustly even when facing data with different levels of illumination, resolution, and frame rate.

### B. Comparison across Time on Selected Videos

Here we compare one of the best performing hand-crafted algorithms, Poh et al. [4], with the DeepPhys model trained in this paper, using a window length of 15 seconds and a sliding step of 1 second. In this setup, we selected the video with the highest accuracy and the one with the lowest for each method. The accuracy values and some frames of these videos are shown in Fig. 3.

We have found that both methods obtain inaccurate results when the videos contain a high amount of movement, when the person blinks with excessive strength, or when he/she does any type of unwitting movement. For both methods, the videos with the best accuracy for heart rate estimation belong to the controlled illumination scenario. On the other hand, the worst results are obtained in videos with uncontrolled illumination. Hand-crafted methods like Poh et al. [4] look for changes in the skin tone at some Regions Of Interest (ROI) inside the face. Due to the unconstrained illumination

conditions at these regions, e.g. one of the cheeks in the images shown in Fig. 3 (right column), the ROI image levels can vary significantly across time according to the varying illumination, degrading the raw rPPG signal and making the accuracy to drop significantly.

In order to illustrate the relevance of the illumination conditions when using hand-crafted rPPG algorithms, we have selected a pair of videos from the testing dataset, one of them captured under controlled illumination conditions and the other with uncontrolled illumination, and we have magnified their color using Eulerian Video Magnification [17], a method that highlights the color changes across the frames. These changes may be caused by the pulse, the external illumination variability, or both. Fig. 4 shows how the color changes due to human pulse are easier to see in the controlled illumination scenario. This explains why hand-crafted rPPG methods such as Poh et al. [4] obtain higher accuracy under these acquisition conditions.

According to the general results shown in Table II, the performance is significantly higher when using DeepPhys compared to the hand-crafted methods. However, there are also cases in which the deep learning algorithm obtains inaccurate results. Taking a look to the MAE obtained for each individual video, it seems that the drop in the accuracy of the deep learning approach is not so intimately related with the illumination conditions, but with other factors like if the users are wearing complements (e.g. glasses) or their

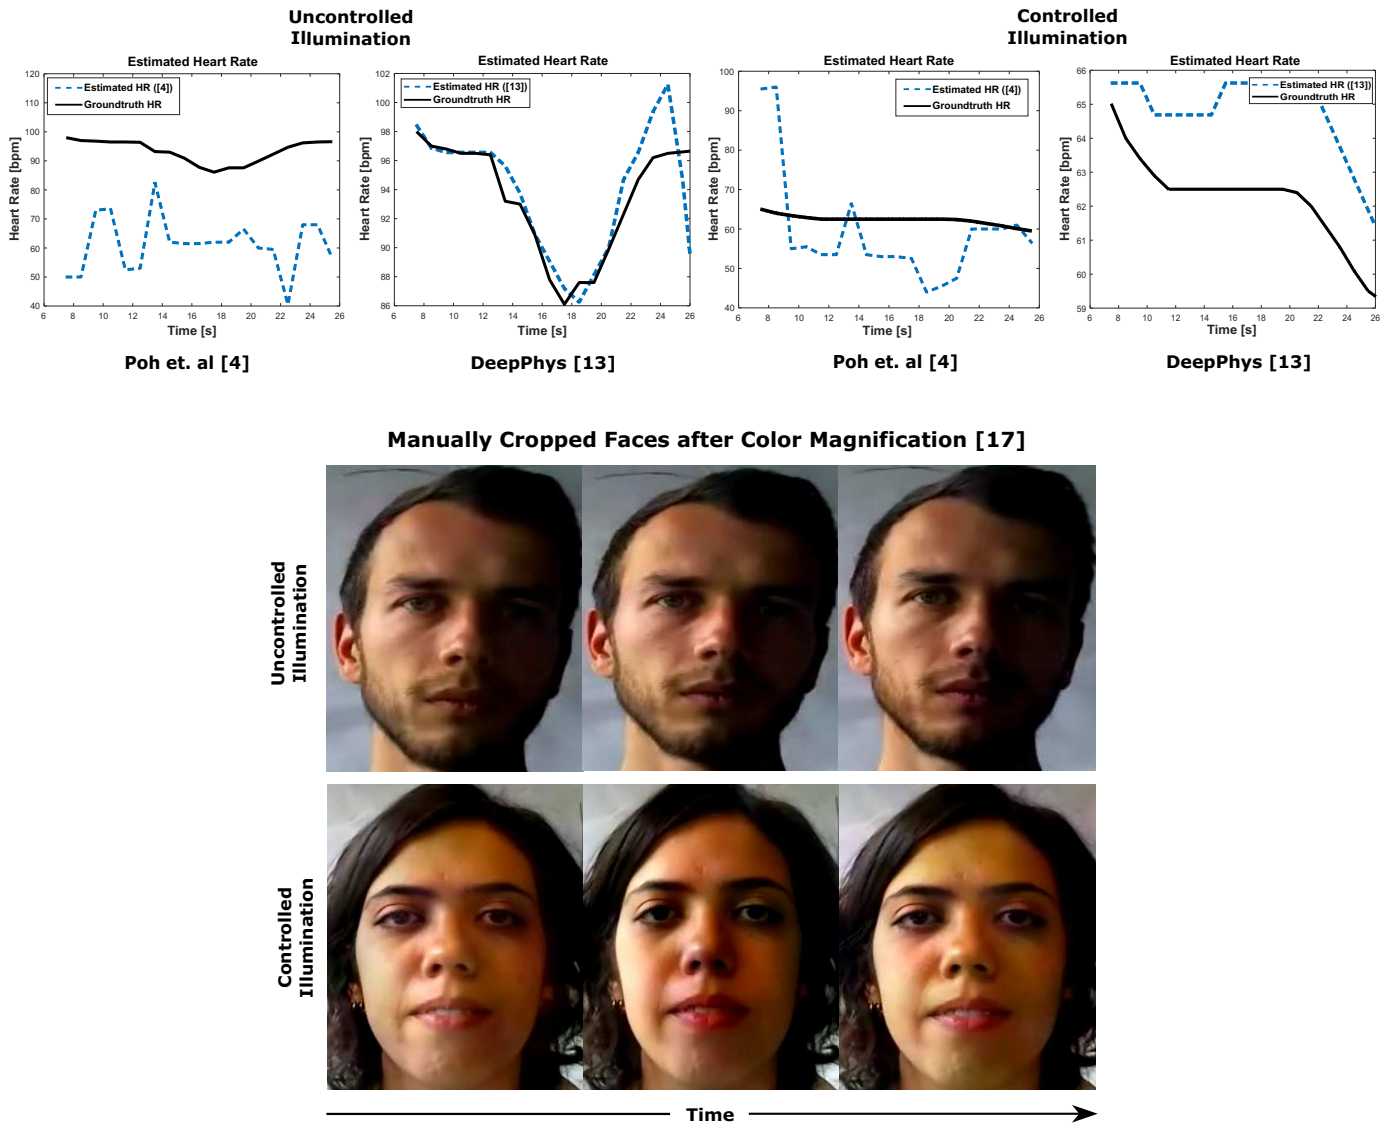


Fig. 4. **Examples of heart rate estimation.** The top row presents the estimation vs. the groundtruth for two videos, one with uncontrolled illumination (that obtained low accuracy with Poh et al. algorithm [4]) and other with controlled illumination (that achieved high accuracy with Poh’s algorithm). For both cases we show the heart rate estimation using [4] and also our DeepPhys [13] model implementation trained with COHFACE. HR is computed for a video length of 15 seconds, with a sliding step of 1 second. The second and third rows show some frames from both videos, whose color has been magnified using [17] for showing how the pulse can be extracted more clearly in the case of controlled illumination acquisition.

skin tone. The training process has made the CNN more robust to illumination variability, since the training set contains the same number of videos with controlled and uncontrolled illumination conditions. This made possible for the network to learn how to extract a precise heart rate estimation in both cases. However, the number of samples of users with dark skin tones or wearing glasses is not so frequent in the training set, so DeepPhys did not learn so accurately how to deal with these conditions. In order to improve the accuracy of DeepPhys it will be interesting to feed the network with a dataset that contains a higher number of these types of samples.

Summarizing, learning-based models have the potential of achieving accuracy values much higher than the ones from hand-crafted approaches, as far as they are properly trained

with enough representative data. These methods are also more robust to external factors such as low illumination or resolution. On the other hand, their main drawback is the need for a high amount of data for training the models, which are not always available.

#### IV. CONCLUSIONS AND FUTURE WORK

In this paper we performed a comparative evaluation of heart rate estimation methods using face videos. This study has taken into account both hand-crafted and learning-based state-of-the-art methods. Results have been provided using the COHFACE database, which is publicly available [15].

For the evaluation we have selected 4 different hand-crafted methods from the most relevant in the literature. Each one of

these methods emphasizes a particular aspect from the vector space generated by the camera sensors. All of these methods have obtained a fairly similar performance.

Under our test scenario conditions, the hand-crafted methods have found a performance barrier around 15 bpm MAE. Human physiology turns out to be complex and highly non-linear and this may be the reason why hand-crafted methods fail to achieve better results.

On the other hand, the suitability of the learning-based models for estimating heart rate measurements through rPPG has been verified. Our DeepPhys implementation has obtained better results than the hand-crafted methods, reaching an error around 5 bpm MAE, similar to the experiments reported by Chen and McDuff in [13], even though in the present work, the COHFACE database contained videos with less controlled illumination than the ones used in the original work, making us think that this type of models could be applied to even more challenging scenarios. The deep neural network is capable of capturing non-linearity thanks to the training process. This result encourages the research in learning-based methods for heart rate detection using face videos, and also for studying other physiological signals that may present a non-linear behavior [19].

For future work it could be interesting to test learning-based methods with input data acquired in less controlled conditions. This could be done by increasing the distance between the camera and the subject, introducing a higher amount of movement in the videos, making illumination conditions even less controlled, etc. The objective is to develop a model capable of accurately estimating a person's heart rate through rPPG from a recording made in the wild [20]. Fine-tuning pre-trained models for adapting them to new scenarios is also an alternative that may be explored, in order to surpass the lack of training data in some cases. Using other parts of the body apart from the face for rPPG is another possibility, e.g. the skin of the torso, the arms, or the legs. This could help to make the heart rate estimation more robust in the cases where the face data is of low quality [21] or directly non available.

## V. ACKNOWLEDGMENTS

This work has been supported by projects: IDEA-FAST (IMI2-2018-15-two-stage-853981), PRIMA (ITN-2019-860315), TRESPASS-ETN (ITN-2019-860813), BIBECA (RTI2018-101248-B-I00 MINECO/FEDER), and edBB (UAM). J.H.O. is supported by a PhD Scholarship from UAM. Portions of the research in this paper used the COHFACE Dataset made available by the Idiap Research Institute, Martigny, Switzerland.

## REFERENCES

- [1] John Allen, "Photoplethysmography and its application in clinical physiological measurement," *Physiological Measurement*, vol. 28, no. 3, 2007.
- [2] Marc Garbey, Nanfei Sun, Arcangelo Merla, and Ioannis Pavlidis, "Contact-free measurement of cardiac pulse based on the analysis of thermal imagery," *IEEE Transactions on Biomedical Engineering*, vol. 54, no. 8, pp. 1418–1426, 2007.
- [3] Wim Verkrusysse, Lars O Svaasand, and J Stuart Nelson, "Remote plethysmographic imaging using ambient light.," *Optics Express*, vol. 16, no. 26, pp. 21434–21445, 2008.
- [4] Ming-Zher Poh, Daniel J McDuff, and Rosalind W Picard, "Advancements in noncontact, multiparameter physiological measurements using a webcam," *IEEE Transactions on Biomedical Engineering*, vol. 58, no. 1, pp. 7–11, 2011.
- [5] Gerard de Haan and Vincent Jeanne, "Robust pulse rate from chrominance-based rPPG," *IEEE Transactions on Biomedical Engineering*, vol. 60, no. 10, pp. 2878–2886, 2013.
- [6] H Emrah Tasli, Amogh Gudi, and Marten den Uyl, "Remote PPG based vital sign measurement using adaptive facial regions," in *Proc. IEEE International Conf. on Image Processing (ICIP)*, 2014, pp. 1410–1414.
- [7] Tong Chen, Peter Yuen, Mark Richardson, Guangyuan Liu, and Zhishun She, "Detection of psychological stress using a hyperspectral imaging technique," *IEEE Transactions on Affective Computing*, vol. 5, no. 4, pp. 391–405, 2014.
- [8] Daniel McDuff, Sarah Gontarek, and Rosalind W Picard, "Improvements in remote cardiopulmonary measurement using a five band digital camera," *IEEE Transactions on Biomedical Engineering*, vol. 61, no. 10, pp. 2593–2601, 2014.
- [9] Jie Chen, Zhuoqing Chang, Qiang Qiu, Xiaobai Li, Guillermo Sapiro, Alex Bronstein, and Matti Pietikäinen, "RealSense = real heart rate: Illumination invariant heart rate estimation from videos," in *Proc. International Conference on Image Processing Theory Tools and Applications (IPTA)*, 2016.
- [10] Xiaobai Li, Jukka Komulainen, Guoying Zhao, Pong-Chi Yuen, and Matti Pietikäinen, "Generalized face anti-spoofing by detecting pulse from face videos," in *Proc. International Conference on Pattern Recognition (ICPR)*, 2016, pp. 4244–4249.
- [11] Wenjin Wang, Albertus C. den Brinker, Sander Stuijk, and Gerard de Haan, "Algorithmic principles of remote PPG," *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 7, pp. 1479–1491, 2017.
- [12] Javier Hernandez-Ortega, Julian Fierrez, Aythami Morales, and Pedro Tome, "Time analysis of pulse-based face anti-spoofing in visible and NIR," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 544–552.
- [13] Weixuan Chen and Daniel McDuff, "DeepPhys: Video-based physiological measurement using convolutional attention networks," in *Proc. of the European Conf. on Computer Vision (ECCV)*, 2018, pp. 349–365.
- [14] Javier Hernandez-Ortega, Shigenori Nagae, Julian Fierrez, and Aythami Morales, "Quality-based pulse estimation from NIR face video with application to driver monitoring," in *Proc. Iberian Conference on Pattern Recognition and Image Analysis*. Springer, 2019, pp. 108–119.
- [15] Guillaume Heusch, Andre Anjos, and Sebastien Marcel, "A reproducible study on remote heart rate measurement," *arXiv preprint arXiv:1709.00962*, 2017.
- [16] Daniel McDuff and Ethan Blackford, "iPhys: An open non-contact imaging-based physiological measurement toolbox," in *Proc. International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2019, pp. 6521–6524.
- [17] Hao-Yu Wu, Michael Rubinstein, Eugene Shih, John Guttag, Frédo Durand, and William T. Freeman, "Eulerian Video Magnification for Revealing Subtle Changes in the World," *ACM Transactions on Graphics*, vol. 31, no. 4, 2012.
- [18] Mohammad Soleymani, Jeroen Lichtenauer, Thierry Pun, and Maja Pantic, "A multimodal database for affect recognition and implicit tagging," *IEEE Transactions on Affective Computing*, vol. 3, no. 1, pp. 42–55, 2012.
- [19] Javier Hernandez-Ortega, Roberto Daza, Aythami Morales, Julian Fierrez, and Javier Ortega-Garcia, "edBB: Biometrics and behavior for assessing remote education," in *Workshop on Artificial Intelligence for Education, AAAI Conference on Artificial Intelligence*, 2020.
- [20] Ester Gonzalez-Sosa, Julian Fierrez, Ruben Vera-Rodriguez, and Fernando Alonso-Fernandez, "Facial soft biometrics for recognition in the wild: Recent works, annotation and COTS evaluation," *IEEE Trans. on Information Forensics and Security*, vol. 13, no. 8, pp. 2001–2014, 2018.
- [21] Fernando Alonso-Fernandez, Julian Fierrez, and Javier Ortega-Garcia, "Quality measures in biometric systems," *IEEE Security & Privacy*, vol. 10, no. 6, pp. 52–62, 2011.