

Face2PPG: An Unsupervised Pipeline for Blood Volume Pulse Extraction From Faces

Constantino Álvarez Casado and Miguel Bordallo López

Abstract—Photoplethysmography (PPG) signals have become a key technology in many fields, such as medicine, well-being, or sports. Our work proposes a set of pipelines to extract remote PPG signals (rPPG) from the face robustly, reliably, and configurably. We identify and evaluate the possible choices in the critical steps of unsupervised rPPG methodologies. We assess a state-of-the-art processing pipeline in six different datasets, incorporating important corrections in the methodology that ensure reproducibility and fair comparisons. In addition, we extend the pipeline by proposing three novel ideas; 1) a new method to stabilize the detected face based on a rigid mesh normalization; 2) a new method to dynamically select the different regions in the face that provide the best raw signals, and 3) a new RGB to rPPG transformation method, called Orthogonal Matrix Image Transformation (OMIT) based on QR decomposition, that increases robustness against compression artifacts. We show that all three changes introduce noticeable improvements in retrieving rPPG signals from faces, obtaining state-of-the-art results compared with unsupervised, non-learning-based methodologies and, in some databases, very close to supervised, learning-based methods. We perform a comparative study to quantify the contribution of each proposed idea. In addition, we depict a series of observations that could help in future implementations.

Index Terms—Remote Photoplethysmography, rPPG, Signal Processing, Pulse rate estimation, Biosignals, Face Analysis.

I. INTRODUCTION

PHOTOPLETHYSMOGRAPHY (PPG) signals have become a key technology in many fields, such as medicine, well-being, or sports. The technique, first introduced by Alrick Hertzman in 1938 [1], utilizes a light source and photodetector to measure blood volume pulse (BVP) through variations in light absorption and reflection in skin tissues [2], [3]. In medicine, PPG analysis is a basic and common tool in healthcare services to monitor vital signs such as heart rate (HR) or oxygen

Manuscript received 22 May 2022; revised 3 May 2023 and 24 June 2023; accepted 17 August 2023. Date of publication 23 August 2023; date of current version 7 November 2023. This work was supported in part by the Academy of Finland 6G Flagship program under Grant 346208, and in part by PROF15 HiDyn program under Grant 326291. (Corresponding author: Constantino Álvarez Casado.)

Constantino Álvarez Casado is with the Center for Machine Vision and Signal Analysis (CMVS), University of Oulu, 90570 Oulu, Finland.

Miguel Bordallo López is with the Center for Machine Vision and Signal Analysis (CMVS), University of Oulu, 90570 Oulu, Finland, and also with the VTT Technical Research Centre of Finland Ltd., 90571 Oulu, Finland (e-mail: miguel.bordallo@oulu.fi).

Digital Object Identifier 10.1109/JBHI.2023.3307942

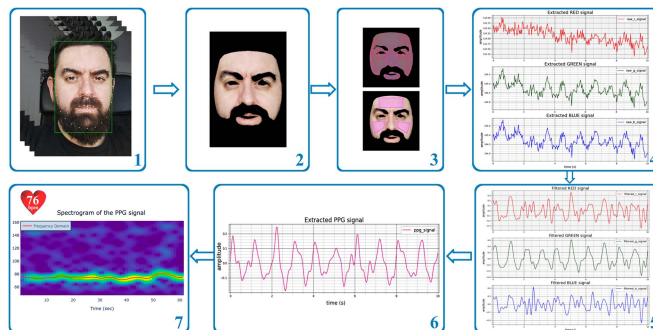


Fig. 1. Typical unsupervised based methodology for remote photoplethysmographic (PPG) imaging using a RGB camera. It comprises several steps: 1) Face detection and alignment 2) Skin segmentation 3) ROI selection 4) Extraction of the raw signals from ROIs 5) Filtered signals 6) RGB to PPG transformation 7) Spectral analysis and post-processing.

saturation (SpO_2) [4]. In well-being, it became increasingly important thanks to the success of wearable devices that analyze sleep disorders [5], cardiovascular diseases [6], or detection of stress and meditation [7]. In sports, PPG analysis became an important tool to improve the intrinsic and extrinsic athletes' performance [8].

Remote PPG (rPPG) imaging is a contactless version of this technology that uses video cameras, usually consumer-grade RGB or near-infrared cameras, and ambient light sources. It works by recording a subject's face or body parts with visible skin areas and analyzing the subtle color variations or motion changes in skin regions [9]. The remote PPG technique allows for non-invasive evaluation and monitoring of users in services, such as healthcare. Hence, the technology could offer significant advantages compared to contact-based devices if it becomes reliable [10].

Current approaches for recovering physiological signals from videos fall into two categories mainly: unsupervised non-learning-based methods, and deep learning approaches [9]. Deep learning-based methods propose end-to-end solutions utilizing training datasets, while unsupervised methods employ computer vision and signal processing in a structured pipeline, as illustrated in Fig. 1.

Most of the unsupervised rPPG methods proposed in the literature focus on recovering PPG signals mostly from static faces, disregarding challenges under real scenarios in real-world applications such as fast face and head movements, extreme light conditions, facial expressions, illumination changes, occlusion, or distance from the camera to the subject. This article focuses on improving the performance of state-of-the-art rPPG unsupervised, non-learning-based methods, emphasizing all system

components. We tackle the improvement of the process by proposing a set of changes across the whole pipeline that result in a noticeable global improvement, performing an extensive evaluation and a framework to recover physiological signals from faces.

A. Contributions

This article depicts different performance problems and challenges to recovering PPG signals from faces reliably. We improve several components of the rPPG pipeline with novel ideas. The main contributions can be summarized as follows:

- We provide a new method to stabilize the movement and facial expression based on a rigid mesh normalization, ensuring that the raw RGB signals are measured from the same facial location regardless of the pose and movement.
- We provide a new method based on statistical and fractal analysis to dynamically select only the facial regions that supply the best raw signals, discarding those with higher noise or prone to artifacts.
- We propose a novel rPPG method to transform the RGB signal into a PPG signal based on QR decomposition, named Orthogonal Matrix Image Transformation (OMIT), which proves to be robust to video compression artifacts.

To prove the usefulness of our approach, we extensively evaluate a set of rPPG methods with four different pipelines across several datasets. Our experiments include modifications to the original evaluation pipeline to increase the fairness and reproducibility of the comparative results.

II. RELATED WORK

In the last few years, rPPG research has progressed from the filtering and simple processing of the variation of the facial skin color to sophisticated multi-step processing pipelines and end-to-end supervised learning methods with dedicated architectures.

A. Unsupervised Methods

Unsupervised non-learning-based methods focus on recovering physiological signals by applying computer vision and signal processing techniques as a system with several steps. These methods focus on obtaining the BVP signal by finding skin areas suitable to extract the raw RGB signals, using face detection, tracking, and segmentation techniques. After that, these methods carefully process these raw RGB signals to separate the physiological signals contained in the subtle variations of the skin color from the rest of the information (motion, illumination changes, or facial expressions, among others) by applying filtering and different ways of combining the RGB signals into an rPPG signal. Most of the studies focus mainly on this transformation component using similar approaches and components for the rest of the process [9].

[11] proposed the first study on extracting remote PPG signals using an inexpensive consumer-grade RGB camera. The study showed how the green channel of the camera contains rich enough, significant information to recover signals such as the heart pulse. [12] proposed the recovery of physiological signals by applying the blind source separation (BSS) technique to remove the noise. Concretely, they used Independent Component

Analysis (ICA) to uncover the independent source signals. Similar to this work, [13] proposed Principal Component Analysis (PCA) to reduce the computational complexity in comparison to Independent Component Analysis with similar accuracy performance. [14] proposed a chrominance-based method (CHROM) to separate the specular reflection component from the diffuse reflection component, which contains pulsatile physiological signals, both reflected from the skin and based on the dichromatic reflection model. [15] define a Blood-Volume Pulse (PBV) vector that contains the signature of the specific blood volume variations in the skin, removing noise and motion artifacts. In the same year, [16] focused on removing the human motions and artifacts from the RGB signals by applying Normalized Least Mean Square (NLMS) adaptive filter. They perform this rectification step by assuming both the face ROI and the background as Lambertian models that share the same light source. [17] proposed a data-driven algorithm by creating a subspace of skin pixels and the computation of the temporal rotation angle of the computed subspace between subsequent frames to extract the heart rate pulse. [18] proposed the CIE Lab color space (LAB) transformation as a more robust color space to extract pulse rate signals due to the high separation between the intensity and chromaticity components, less sensitive to human body movements. The study also demonstrates that the a channel has a better signal-to-noise ratio (SNR) than the green channel in RGB color space. [19] proposed a new plane-orthogonal-to-skin (POS) algorithm that finds pulsatile signals in an RGB normalized space orthogonal to the skin tone. [20] proposed the Local Group Invariance (LGI) method, a stochastic representation of the pulse signal based on a model that leverages the local invariance of the heart rate as a quasi-periodical process dynamics and obtained by recursive inference to remove extrinsic factors such as head motion and lightness variations. [21] proposed an unsupervised method with an emphasis on motion suppression and novel filtering based on the head orientations (FaceRPPG). Recently, deep learning methods have been developed that do not rely on reference signals, which could be considered unsupervised [22], [23]. However, they remain specifically tailored to each dataset and its unique characteristics.

Unsupervised non-learning-based methods present significant advantages, as they do not rely on specific training data, enabling better generalization across different datasets and measurement setups. These methods focus on measuring the BVP signal as it manifests in various facial regions. However, the performance of unsupervised methods can be influenced by several factors, such as sensitivity to noise and artifacts, dependency on user-defined parameters, and limited adaptability to varying skin types and lighting conditions.

Unsupervised non-learning-based method's efforts have focused on finding suitable ways of transforming noisy RGB signals into reliable PPGs. In contrast, the impact of other system components, such as face detection and tracking, has been mainly disregarded. Our contribution addresses the unsupervised rPPG process as a system with multiple components that can be improved separately.

B. Supervised Methods

Deep Neural Networks (DL) and especially *Convolutional Neural Networks* (CNN) approaches have gained attention and become popular tools in computer vision and signal processing tasks, including healthcare-related tasks. Before the advent

of deep-learning based methods, there were two primary approaches to estimate heart rate using machine learning methods, including support-vector regression [24] and adaptive hidden Markov models [25]. Since 2018, supervised deep learning-based methods to compute HR or other vital signs started to arise increasingly in the literature. Among the most relevant learning-based remote PPG methods, [26] proposed a two-step convolutional neural network to estimate a heart rate value from a sequence of facial images. HR-CNN is a trained end-to-end network composed of two components, an Extractor and an HR Estimator. The same year, [27] proposed DeepPhys, another end-to-end solution based on a deep convolutional network that estimates HR and breathing rate (BR). The approach performs a motion analysis based on attention mechanisms and a skin reflection model using appearance information to extract the physiological signals. [28] proposed RhythmNet, an end-to-end solution based on spatial-temporal mapping to represent the HR signals in videos. The approach also exploits temporal relationships of adjacent HR estimations to perform continuous heart rate measurements. The same year, [29] proposed a two-stage end-to-end solution. The first part of the network, named STVEN, is a Spatio-Temporal Video Enhancement Network to improve the quality of highly compressed videos. The second part of the approach, called rPPGNet, is the 3D-CNN network that recovers the rPPG signals from the enhanced videos. The authors claim that the proposed rPPGNet produces rich rPPG signals with curve shapes and peak locations. [30] further proposed another end-to-end approach for remote HR measurement based on Neural Architecture Search (NAS). AutoHR is comprised of three ideas: a first stage that discovers the best topology network to extract the physiological signals based on Temporal Difference Convolution (TDC); a hybrid loss function based on temporal and frequency constraints; and Spatio-temporal data augmentation strategies to improve the learning stage. The same year, [31] proposed a transductive meta-learner based on an LSTM estimator and Synthetic Gradient Generator that adjusts network weights in a self-supervised manner. Recently, [32] proposed a two-stage hybrid method called PulseGAN. It starts with the unsupervised extraction of noisy PPG signals using CHROM as a first stage, followed by a generative adversarial network (GAN) that generates realistic rPPG pulse signals from the signals recovered in the first stage. In 2021, [33] proposed a novel denoising-rPPG method called AND-rPPG, based on the utilization of Action Units (AUs) and Temporal Convolutional Networks (TCNs) for denoising temporal signals to mitigate facial expression noises effectively.

Supervised rPPG methods, predominantly employing deep learning techniques, boast remarkable accuracy by utilizing end-to-end solutions. With minimal intermediate steps required beyond essential preprocessing, these methods learn to recognize facial noise patterns in relation to reference contact-based PPG ground-truth signals obtained from the finger [34]. However, this approach results in a black-box model that recovers rPPGs from video frames without understanding the underlying mechanisms, posing significant challenges in critical domains such as healthcare [35], e.g. for cardiovascular conditions. Their reliance on training data limits their generalizability to new scenarios, subjects, or recording conditions and raises concerns regarding data acquisition and cost. Moreover, supervised methods may need substantial computational resources.

III. UNSUPERVISED BLOOD VOLUME PULSE EXTRACTION METHODOLOGY

Our work uses a standard methodology in unsupervised rPPG approaches to extract the blood volume changes from facial videos and derive essential parameters. We follow a modular pipeline with several components, as depicted in Fig. 1. The pipeline is roughly divided into three big blocks: the selection of measuring regions of the face, the extraction of rPPG biosignals from natural variations in color or texture, and the computation of the heart rate or other parameters using the extracted signals.

The pipeline comprises of 8 main modules sequentially connected:

- 1) *Database interface*: Provides interfaces to read data from various public databases, including videos, images, and reference signals.
- 2) *Face detection and alignment*: An initial step for detecting and aligning faces in each frame. This process obtains face coordinates and facial points, which are crucial for selecting regions of interest, patches, or segmentation coordinates.
- 3) *ROI selection*: step to select the regions of interest of the face based on a color skin segmentation or selection of patches based either on face location or the coordinates of the landmarks.
- 4) *RGB extraction*: Retrieves raw signal from a window of multiple RGB frames, computing the mean value of pixels within the mask or patches selected earlier.
- 5) *Pre-processing*: Block mainly to filters raw RGB signal to focus on the band of interest. Filtering is essential for accurate BVP signal recovery, as remote PPG signals often exhibit trends and noise. The heart-related signal lies between 0.75 and 4.0 Hz.
- 6) *RGB to PPG transformation (rPPG)*: crucial and core step in remote PPG methods as it converts skin color variations into physiological signals. This step involves transforming the RGB signal into a PPG signal using a transformation method that combines the RGB channels to generate a pulse signal.
- 7) *Frequency analysis*: spectrum analysis of the rPPG and reference ground-truth (BVP or ECG) signals to estimate the heart rate.
- 8) *Evaluation*: it includes an error and statistical analysis to compare the estimated heart rate from the rPPG and the estimated heart rate from the ground-truth signals.

A. Baseline Pipeline

In our work, we start from an open-source framework for the evaluation of remote PPG methods, implemented in Python and called *pyVHR* (short for Python tool for Virtual Heart Rate) [36]. This framework includes an extensible interface to integrate several datasets, multiple methods, choices for each processing step, and extensive assessment and visualization tools. We use version 0.0.4 of the PyVHR framework, which we will refer to as *Baseline* pipeline onward.




Fig. 2. Facial landmark detection using DAN model under extreme head poses and frontal faces.

IV. FACE2PPG PIPELINES

The *Baseline* pipeline presents a few shortcomings that might result in inaccurate or unfair assessments. We incrementally improve it by introducing several changes in multiple steps, and propose three new versions that we name *Improved*, *Normalized* and *Multi-region* pipelines.

These pipelines focus on handling the face in unconstrained conditions since it is one of the critical parts of extracting remote photoplethysmograms. It has to be noted that most of the unsupervised approaches have focused mainly on developing RGB to PPG conversion methods (sometimes called just rPPG methods), but they have not paid much attention to other steps of the pipeline. This article emphasizes the importance of every step when extracting and evaluating remote PPG signals from faces. Moreover, we have devised a novel method, Orthogonal Matrix Image Transformation (OMIT), which employs QR decomposition to convert raw RGB signals into BVP signals.

A. Improved Pipeline

To mitigate its shortcomings, we modified the *Baseline* pipeline to incorporate a few minor changes that increase reproducibility and enable a fairer comparison of different methods. We name this modified version as *Improved* pipeline. We enumerate and describe these changes as follows:

Face detection: The *Baseline* pipeline includes two well-known face detectors: one based on convolutional neural networks known as MTCNN [37] and a *Dlib* implementation based on Histogram of Oriented Gradients (HOG) features [38]. We use instead a new deep learning-based face detection method based on a Single Shot Multibox Detection network (SSD) [39], implemented in *OpenCV* library. This face detector outperforms the *Baseline* pipeline detectors in terms of accuracy, size of the models, and computational speed [40].

Face alignment: The *Baseline* pipeline includes two well-known face landmark detectors, the MTCNN detector [37] that computes 5 landmark points in the face (eyes, nose and mouth corners) and the *Dlib* implementation of the ERT method [41], [38]. We use instead a deep learning approach named DAN (Deep Alignment Network) [42] which gives exceptional performance in terms of accuracy, even in challenging conditions [43] as shown in Fig. 2. For a faster, real-time face alignment, we have added a more accurate, faster, and smoother model for the ERT *Dlib* face landmarks detector [43]. These models both infer 68 landmark points defined by the Multi-PIE landmark scheme.

Filtering: The *Baseline* pipeline only considers a pre-filtering scheme before the RGB to PPG transformation. The pipeline offers three types of filters: detrending (Scipy or Tarrainen methods), bandpass filtering (FIR filter with Hamming window and Butterworth IIR filter), and a Moving average filter




Fig. 3. Delay between reference BVP and remote PPG signals from proposed pipelines, induced by factors like filtering, blood perfusion differences, or camera distance.

(MA) that removes various base noises and motion artifacts of the signals. We have added the possibility of using Kaiser windows when applying FIR filtering. A Kaiser-Bessel window maximizes the energy concentration in the main lobe, and it is highly recommended to filter biosignals [44]. In addition, we have introduced the possibility of applying also post-filtering, performed after the RGB to PPG conversion, since the literature suggests that some conversion methods perform better this way [45].

RGB to PPG transformation (rPPG): The *Baseline* pipeline includes several reference methods such as POS [19], CHROM [14], GREEN [11], PCA [13], ICA [12], SSR [17], LGI [20] and PVB [15]. We have added one method based on selecting the chroma channel *a* after applying a CIE Lab color space transformation. CIE Lab separates the lightness information (channel L) from the chroma information (channels *a* and *b*). The chrominance components have a more significant dynamic range than the red, green, and blue channels in RGB color space [18]. In addition, it correlates with skin color and related parameters and describes better the subtle changes occurring in them [46].

Spectral analysis: In the original framework, the ground-truth pulse rate is estimated using Short Time Fourier Transform (STFT) when the ground-truth signal is a BVP signal, and R-Peak detection and RR interval analysis when the ground-truth signal is an ECG signal. In the *Baseline* pipeline, the recovered PPG signals are processed instead of using Welch's spectral density estimation. This mismatch introduces the possibility of unfair evaluation. We modified the pipeline, so the ground-truth BVP signal and the rPPG signal are processed using the same spectral analysis algorithm and similar parameters such as the overlap or FFT length.

Evaluation: It can be expected that the reference BVP (PPG) signals taken in the finger and the rPPG signals extracted from the face are not perfectly synchronized or show the same dynamic range. We show an example in the Fig. 3. We can also observe asynchrony in the heart rate estimation, as shown in Fig. 4.

Technical and physiological factors can produce time shifts and morphological differences in the signals. These effects can be attributed to the distance between the optical sensors, the contact force effect of the finger PPG oximeter [47], different filtering parameters [48], individual variations among subjects [49], variability in the measurement site [50], and even blood perfusion differences in different body regions [51], among others. We mitigate some of the effects caused by




Fig. 4. Estimated HR from the reference BVP signal and the extracted facial rPPG signal using the *Baseline* pipeline. We can observe the asynchrony due to the different signal sources.

comparing fundamentally different signals by adding a new parameter in the dataset interface that aligns both signals in terms of time and dynamic range, resulting in a fairer estimation of the error between ground-truth HR estimation and rPPG HR estimation. The parameter has been adjusted for each dataset globally using a mostly empirical approach due to the lack of detailed information on the measurement setups. This approach ensures the consistency of signal alignment while maintaining the overall integrity of the data.

B. Normalized Pipeline

Our *Normalized* pipeline introduces two significant changes compared to the previous ones, a segmentation approach based on geometric normalization, and a novel RGB to PPG transformation method, robust to compression artifacts.

1) *Geometric Segmentation and Normalization*: One of the critical steps in the non-contact PPG extraction is the process related to skin segmentation, since it is the source to recover the desired physiological signals. Most of the unsupervised methods and pipelines rely on simple thresholding of different color spaces for skin color segmentation [36], from inefficient fixed RGB segmentation to adaptive HSV segmentation. These pixel-level techniques suffer from generalization due to the variability of skin tones, skin paint, illumination changes, and complex backgrounds. It is not easy to define clear boundaries between the skin and non-skin pixels, mainly due to the variability of the facial regions measured across frames of a single video. Framewise skin segmentation based on neural networks (e.g., uNET) suffers similar problems, sometimes caused by the small number of annotated facial skin masks, resulting in underfitted models [52].

We propose using a geometrical segmentation scheme that uses fiducial landmark points detected in the face. Although some interframe jittering due to the landmark variability remains and produces changes in the skin mask across the video frames, this is noticeably lower than using skin color segmentation. To perform this segmentation, we have extended the set of landmark points from 68 to 85 landmark points by interpolation and created a fixed facial mesh composed of 131 triangles and fix their coordinates as a typical frontal face, as shown in Fig. 5.

Our segmentation approach normalizes the face in each frame by mapping every triangle in the current detected face to the triangles in the normalized shape as shown in Fig. 5. This approach generates a spatio-temporal matrix of normalized faces that ensures that we measure the signals in the same facial




Fig. 5. Face normalization process. Left to right: detected landmark points, fixed triangle face mesh, and normalized face after mapping triangles to fixed coordinates.

regions consistently across frames, regardless of the pose and movement.

C. Multi-Region Pipeline

Extracting only one signal from the whole face or skin mask can result in a very noisy signal. However, the previously described skin segmentation produces a set of facial regions that can be analyzed separately. Due to partial occlusions, extreme head poses, illumination variations, or shades, among others, some of these regions might present very noisy signals with low dynamic ranges. Previous methods have proposed to select fixed patches in the face, where a priori, the blood perfusion should be more observable (e.g., forehead and cheeks). This approach works relatively well when the videos show quasi-static individuals in fixed environments but fails when presented with fast movements or strong face rotations. To mitigate the impact of these challenges, we propose to modify the pipeline by introducing a dedicated block that automatically and dynamically selects those regions that contain the raw signals with higher quality. We name the resulting framework as *Multi-Region Pipeline*.

1) *Dynamic Multi-Region Selection*: We propose a novel Dynamic Multi-Region Selection (DMRS) method to select the best facial regions dynamically. This approach extracts signals in a fixed set of facial regions and statistically analyzes their quality to choose whether each one of them should contribute to the final rPPG signal or if it should be discarded.

The DMRS process starts just after obtaining a segmented and normalized face in the previous block of the pipeline by dividing the normalized face into a matrix of $n \times n$ rectangles (regions of interest) that contains a spatio-temporal representation of the face, as depicted in Fig. 6. Each area in the grid represents a signal in a sequence of frames.

Next, the process continues the analysis by computing several statistical parameters on each candidate region and the global face. These statistical parameters are computed using windows of t seconds, both in time and frequency domains. We extract the mean, standard deviation, variance, signal-to-noise ratio (SNR), Katz Fractal dimension (KFD), number of zero-crossings (Z_c), sample entropy, detrended fluctuation analysis (DFA), and the energy in terms of local power spectral density (PSD). The dynamic selection is loosely based on fractal analysis, which portrays the scale of the randomness or how unpredictable a stochastic process is [53].

The first step is an initial pruning that removes those regions that do not contain valuable information. We check the variance changes of every candidate rectangle along the time t and discard those that have zero variance.




Fig. 6. Comparison of candidate regions of interest in a frontal face under good lighting conditions (top row) and a non-frontal face under suboptimal lighting conditions (bottom row). The system selects 32 regions in the frontal face and 20 regions in the non-frontal face. Grey regions are discarded.

Then, the regions are discarded based on the thresholding of Katz’s Fractal Dimension (KFD). KFD computes the fractal dimension (FD) of a signal based on the morphology, measuring the degree of irregularity and the sharpness of the waveform [53]. KFD gives an index D_{KFD} for characterizing the complexity of the signal. This value is computed as shown in the following equation:

$$KFD = \frac{\log_{10}(L/a)}{\log_{10}(d/a)} = \frac{\log_{10}(n)}{\log_{10}(d/L) + \log_{10}(n)} \quad (1)$$

where L is the total length of the PPG time series, a is the average of the Euclidean distances between the successive points of the sample, d is the Euclidean distance between the first point in the series and the point that provides the maximum distance with respect to the first point, and n is L/a . We calculate then the relative D_{KFD} value by dividing the KFD index of a specific facial region by the global KFD index derived from the entire face. This relative value plays an important role in identifying the regions that contribute meaningful information to the final rPPG signal while addressing the inherent challenges posed by facial regions, such as occlusions, low light conditions, blur, and other factors that could adversely affect signal quality and the complexity of the signal. By selecting regions with a relative D_{KFD} value of 0.85 or greater, we ensure the inclusion of regions that exhibit complexity levels comparable to the global PPG signal. This approach is based on the assumption that the primary source of complexity in the global rPPG signal stems from the heart, and regions with similar complexity levels are more likely to provide valuable information for rPPG analysis, despite the presence of factors that could potentially compromise the signal quality. Through this selection process, we effectively discard regions that introduce noise and artifacts in the final rPPG signal. This method provides a good balance between discarding low-quality regions and retaining meaningful information.

The analysis continues by using Detrended Fluctuation Analysis (DFA), a statistical method widely used to detect intrinsic self-similarity in non-stationary time series, especially in fractal signals. DFA is a modified root mean square analysis of a random walk, designed to compute long and short-range non-uniform

correlations in stochastic processes [54]. The method tells if each region rPPG signal shows the expected correlation with the signal from the global face signal and if they are very noisy or contain artifacts of extrinsic trends [55]. The DFA exponent α is interpreted as an estimation of the Hurst parameter, and it is calculated as the slope of a straight line fit to the log-log graph from the fluctuation function. If $\alpha = 0.5$, the time series is uncorrelated. If $0.5 < \alpha < 1$ then there are positive correlations in the time series. If $\alpha < 0.5$ then the time series is anti-correlated. We discard those regions as uncorrelated and negatively correlated.

Upon completing the analysis of each facial region in a frame sequence (window of t seconds), we obtain a set of valid regions (2 to 32) for the energy-based final selection. If more than r_{\max} regions pass the analysis, we choose the top r_{\max} regions with the highest energy, reducing noise and improving the signal-to-noise ratio [56], [57]. If there are fewer than r_{\max} but more than one valid regions, all are selected. If only one or none qualify, we revert to selecting the best r_{\max} regions by energy content from the initial candidates. This strategy optimizes rPPG signal quality across scenarios. Lastly, the rPPG signals from chosen regions are combined by summing in the time domain, generating the final rPPG signal for heart rate computation during spectral analysis.

V. ORTHOGONAL MATRIX IMAGE TRANSFORMATION

We introduce a novel, robust, and efficient RGB to PPG transformation method called Orthogonal Matrix Image Transformation (OMIT), that we integrated in the previous proposed pipelines.

OMIT leverages matrix decomposition techniques, generating an orthogonal matrix with linearly uncorrelated RGB components. It employs reduced QR factorization [58] and Householder Reflections [59] for finding linear least-squares solutions in the RGB space. The thin QR factorization offers memory efficiency and computational speed, especially for tall and skinny matrices [58], while the Householder Orthogonalization Algorithm provides better numerical stability, computational efficiency, and conditioning compared to the Gram-Schmidt process [58], [59]. These advantages make OMIT well-suited for handling noisy data matrices and extracting rPPG signals with greater accuracy and efficiency.

The mathematical foundation of OMIT is based on the QR decomposition as shown in (2):

$$A = QR \quad (2)$$

where $A \in \mathbb{R}^{n \times 3}$ represents the input RGB matrix, $Q \in \mathbb{R}^{n \times 3}$ denotes the orthonormal basis for the column space of A , and $R \in \mathbb{R}^{3 \times 3}$ is an upper triangular matrix that contains the coefficients to express the columns of A as linear combination of the basis vectors in Q . We then use the orthogonal matrix Q to compute a projection matrix that allows us to extract the BVP signal from the input matrix A .

The OMIT method is composed of the following key steps, illustrated in Fig. 7:

- 1) *Reduced QR decomposition using Householder Reflections*: Compute the thin QR decomposition of the input RGB matrix [58], [60]. After k iterations (in our case, $k = 3$), the product of the Householder Reflectors (H_i) matrices forms the semi-orthogonal matrix Q , and the transformed A becomes the upper triangular matrix R .

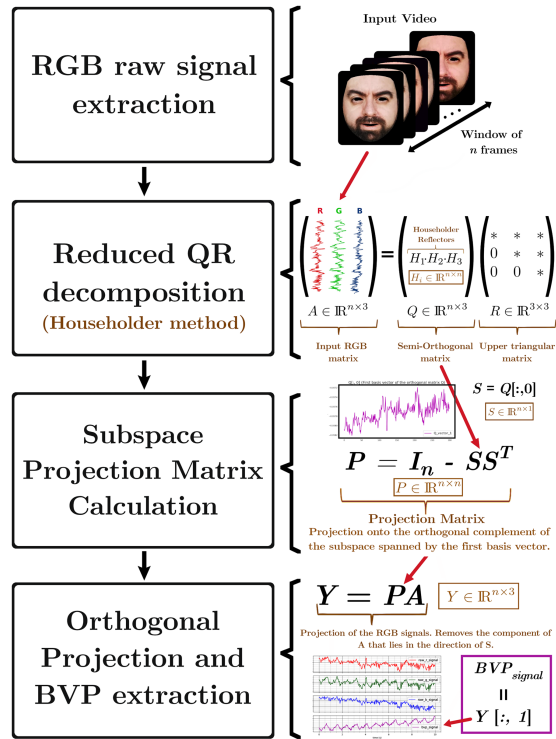


Fig. 7. Process steps to convert a RGB signal to BVP signal using OMIT method.

The first column of Q captures the most significant variations in the input data.

- 2) *Subspace projection matrix calculation:* The first column of Q , denoted as S , is used to compute the projection matrix P , which is calculated as $P = I_n - SS^T$. This matrix is designed to project the input data onto a subspace orthogonal to S , effectively removing the contributions associated with the dominant variations in the input data.
- 3) *Orthogonal Projection and BVP extraction:* The input data (RGB matrix) is projected onto a subspace orthogonal to S using the calculated projection matrix P , represented by the equation $Y = PA$. This step separates the BVP signal from the raw RGB data while suppressing noise and artifacts. The BVP signal is extracted from the second column of the Y matrix after removing the dominant variations.

QR decomposition has been extensively applied across various domains such as communications, signal processing, image processing, and machine learning to address challenges associated with corrupted input data matrices due to noise or artifacts [61], [62].

QR decomposition has shown superior performance over alternative techniques such as Principal Component Analysis (PCA), Eigenvalue Decomposition (EVD), and Singular Value Decomposition (SVD) in terms of mathematical stability, computational efficiency, and robustness [63]. OMIT utilizes these benefits to produce an orthogonal matrix with linearly uncorrelated components, effectively isolating the rPPG signal from RGB data. This allows for more accurate and efficient blood pulse signal extraction, enhancing rPPG signal processing and analysis.

VI. BENCHMARK DATASETS AND EVALUATION METRICS

We evaluate our methodology using an extensive assessment based on the literature [36] and six publicly available datasets: PURE [64], COHFACE [65], LGI-PPGI-Face-Video-Database [20], UBFC-RPPG Video dataset [66], and MAHNOB-HCI [67]. These datasets consist of videos captured under various conditions and setups, as well as reference physiological data recorded using medical-grade equipment such as fingertip pulse oximeters and ECG sensors.

PURE contains 10 subjects performing controlled head motions with 60 uncompressed sequences and uses a fingertip pulse oximeter as ground truth. **COHFACE** includes 160 videos of 40 subjects synchronized with heart rate and breathing rate and uses medical-grade equipment for reference data. **LGI-PPGI-Face-Video-Database** comprises four different scenarios with 25 subjects, using an uncompressed format and a fingertip pulse oximeter as ground truth. **UBFC-RPPG** Video dataset has two sub-datasets: UBFC1 with 8 uncompressed videos and UBFC2 with 42 uncompressed videos, presenting diverse ethnicities and facial skin tones, and uses a fingertip pulse oximeter for reference data. **MAHNOB-HCI** contains 27 participants with 527 compressed facial videos and corresponding physiological signals, using ECG as ground truth. For our evaluation, we used a smaller subset of 36 videos from the MAHNOB-HCI database to ensure a fair comparison with the results in [36] and the full dataset of 527 videos for comparing our pipeline with other state-of-the-art approaches.

The evaluation of the datasets is done by comparing the estimations of the heart rates of both the extracted rPPG signal and the reference ECG or PPG signal. The evaluation includes both error and statistical analysis. We use three standard metrics that measure the discrepancy between our predicted heart rate $\hat{h}(t)$ and the reference heart rate $h(t)$. The standard metrics used to compute it are Mean Absolute Error (MAE), Root-Mean-Square Error (RMSE), and Pearson Correlation Coefficient (PCC) of the heart-rate envelope. Our primary objective is to advance unsupervised rPPG extraction techniques in challenging conditions, rather than pursuing waveform similarity with ground truth PPG signals, as is the case with deep learning-based methods. Given the anatomical differences in blood perfusion waveforms between the face and finger [68], we employ PCC between heart rate envelopes as a more suitable evaluation metric, instead of direct waveform comparisons.

A. Reference Data, Ground-Truth and Evaluation Protocol

The most common source of BVP reference data in the datasets is PPG data from contact-based pulse oximeters. In most of the cases, these data is already filtered and does not require further preprocessing. In order to extract heart rate or other HRV parameters, the reference signals are processed using spectral analysis.

In the evaluation, we compute the error by comparing the heart rate and HRV parameters extracted from the reference (ground-truth) signals and the recovered PPG signal. Although direct comparison of signals (e.g. morphology) would be also possible, the fundamental differences between the extracted rPPG and the reference signals in terms of delay and scale due to different body measurement points and diverse collection devices make this comparison not very meaningful [68].




Fig. 8. Reference PPG signal from a fingertip contact-based pulse oximeter, showing a ≈ 2 -second error, likely due to finger movement and signal tracking loss.

We have observed that the reference data offered in the datasets is not completely free of problems. For example, Fig. 8, shows an example reference signal with a gap of approximately 2 seconds. These issues, caused by small deficiencies in data collection can lead to unfair disagreements in terms of error, especially for unsupervised methods.

VII. EXPERIMENTAL RESULTS

We evaluate and analyse the proposed methodologies and pipelines to extract remote photoplethysmography signals from all benchmark databases. We compare the results across different improved processing pipelines and compare them with the state of the art for both supervised and unsupervised methods. The experiments are performed using a computer that includes an AMD Ryzen(TM) 3700X 8-core processor at 3.6 GHz.

A. Hyperparameters and Configuration

Our framework is based on separate configuration files, in a similar manner as other frameworks [36]. These files contain the parameters that govern the pipeline and its components. In our experiments, we set the values for of each pipeline component as follows: **Face detection** uses a DNN OpenCV face detector with the default Tensorflow model. **Face alignment** uses the DAN algorithm with one of the default models provided by the authors (*DAN-Menpo.npz*) [42]. Real-time configurations can use a modified ERT model [43]. **DMRS** uses a grid matrix of $n = 9 \times 9$. For D_{KFD} , we have selected the signals with a threshold of more than 0.85, and for DFA we set an α threshold between 0.75 and 1.0. We set the maximum number valid of regions $r_{max} = 32$. **Filtering** is performed using FIR filters with Kaiser windows, with the parameter $\beta = 25$. The filters use a bandpass configuration between 0.75 and 4 Hz (corresponding to 45–240 bpm). **Signal Windowing** uses sliding windows of 10 seconds and 1 s steps (9 seconds overlap).

B. Quantitative Results

We provide an extensive evaluation of our three proposed pipelines and compare them with those of the baseline [36] with the standard configuration. We obtain results in six datasets. All datasets are comprised by videos with VGA resolution, but in two of them are heavily compressed (MAHNOB and COHFACE).

We measure the performance by computing the average of the MAE, the standard deviation of the MAE, and the median of



Fig. 9. Comparison of HR estimation using four rPPG methods with the Multi-region pipeline. Estimated heart rate (red) is extracted from the face, and the contact-based reference PPG signal (blue) is computed on a single PURE dataset video.

the Pearson Correlation Coefficient of the envelope of the heart rate. We evaluate the pipelines using ten rPPG methods (RGB to PPG signal conversion methods), including our proposed OMIT conversion. The results detail the impact of the improvements in each pipeline as shown in Table I.

The Multi-region pipeline, with our proposed improvements, achieves the best results across all six datasets, improving MAE, error standard deviation, and Pearson Correlation Coefficient of the heart rate envelope.

Analyzing different rPPG conversion methods, CHROM and POS perform best in uncompressed databases across all pipelines, with OMIT closely following. OMIT works well in highly compressed pipelines, obtaining the best results for the challenging MAHNOB dataset.

Comparing results across datasets, the mean average error varies significantly depending on the data nature. Good quality and static datasets like UBFC or PURE have an error below 2 bpm, with minor differences across videos. For the LGI-PPGI dataset with natural movement, the best average error reaches nearly 4 bpm, with a reasonably high standard deviation. Worst results correspond to lower resolution datasets like MAHNOB and COHFACE, with average errors between 8 and 12 bpm. Heavy video compression and low illumination can cause low SNR and loss of signal subtleties [69].

Although some variations across methods, dataset, and pipelines exist, it is possible to conclude that the modification introduced in the pipelines shows a consistent improvement, although those databases with faces mostly still in front of the camera (PURE and UBFC), show only modest improvements when compared with those achieved in complicated datasets (LGI-PPGI, COHFACE, MAHNOB).

We argue that the proposed multi-region pipeline shows better performance when presented with videos collected in a more natural environment, especially when they show varied facial expressions, head movements, or illumination changes. To illustrate this, we evaluate results on the LGI-PPGI dataset across four scenarios: resting (still subjects with no head or facial movements), rotation (head motions and rotations), talking (street video conference with sudden head motions and low dynamic range), and gym (exercise on a static bicycle), using our proposed OMIT method for RGB to PPG conversion. Baseline and Multi-region pipelines perform similarly in resting and rotation scenarios. However, in challenging scenarios like talking and gym, errors significantly reduce and the Pearson Correlation

TABLE I
ERROR COMPARISON BETWEEN THE BASELINE PIPELINE AND OUR THREE PROPOSED IMPROVED PIPELINES

Pipeline	PPG Method	Databases											
		LGI-PPGI		COHFACE		MAHNOB		PURE		UBFC1		UBFC2	
		MAE ± SD	PCC	MAE ± SD	PCC	MAE ± SD	PCC	MAE ± SD	PCC	MAE ± SD	PCC	MAE ± SD	PCC
Baseline (+ PPG Method)	GREEN (2008)	16.3 ± 11.4	0.23	12.3 ± 7.2	0.05	15.3 ± 7.5	0.01	5.8 ± 12.9	0.64	15.9 ± 10.4	-0.04	8.1 ± 11.7	0.60
	ICA (2010)	11.2 ± 11.4	0.22	11.1 ± 5.8	0.20	14.9 ± 6.2	0.17	4.4 ± 5.2	0.36	5.8 ± 6.0	0.44	10.4 ± 7.2	0.41
	PCA (2011)	10.9 ± 12.1	0.22	10.4 ± 5.7	0.27	11.0 ± 5.4	0.17	1.9 ± 1.7	0.77	4.5 ± 1.8	0.43	5.2 ± 3.9	0.58
	CHROM (2013)	11.1 ± 14.2	0.39	12.4 ± 7.0	0.03	18.6 ± 9.0	0.06	1.6 ± 2.0	0.83	2.2 ± 0.8	0.71	3.1 ± 2.5	0.79
	PBV (2014)	17.4 ± 13.1	0.03	13.2 ± 6.5	0.05	18.6 ± 6.5	-0.02	3.6 ± 3.7	0.45	8.9 ± 6.0	0.30	5.5 ± 7.0	0.63
	2SR (2016)	10.1 ± 8.9	0.31	11.8 ± 7.5	0.02	17.5 ± 5.3	0.05	3.7 ± 7.1	0.73	4.1 ± 3.8	0.64	5.7 ± 5.6	0.40
	LAB (2016)	10.9 ± 10.8	0.35	11.2 ± 5.6	0.03	20.5 ± 10.3	0.05	2.4 ± 3.5	0.64	7.5 ± 7.8	0.39	12.5 ± 13.2	0.39
	POS (2017)	11.8 ± 14.4	0.38	11.9 ± 7.1	0.05	19.4 ± 5.5	0.02	1.9 ± 3.1	0.85	1.8 ± 0.4	0.87	1.9 ± 1.4	0.91
	LGI (2018)	10.3 ± 13.2	0.42	12.6 ± 7.1	-0.01	19.0 ± 5.5	0.05	3.2 ± 3.5	0.59	2.4 ± 1.0	0.70	8.7 ± 6.3	0.40
	OMIT (2022)	9.5 ± 11.1	0.48	12.4 ± 6.6	0.04	19.9 ± 9.0	0.13	2.3 ± 4.9	0.65	5.9 ± 7.1	0.89	5.8 ± 8.2	0.82
Improved (+ PPG Method)	GREEN (2008)	9.3 ± 7.9	0.27	9.6 ± 5.7	0.01	12.2 ± 5.8	0.16	4.8 ± 8.5	0.65	8.7 ± 10.2	0.68	5.4 ± 5.9	0.81
	ICA (2010)	7.7 ± 9.3	0.30	8.8 ± 4.8	0.07	17.0 ± 7.0	0.06	1.9 ± 3.9	0.86	6.3 ± 8.2	0.77	6.6 ± 8.0	0.66
	PCA (2011)	7.8 ± 9.9	0.35	8.6 ± 4.8	0.04	17.2 ± 8.9	0.10	1.9 ± 4.3	0.86	5.5 ± 6.5	0.59	5.0 ± 6.7	0.95
	CHROM (2013)	4.8 ± 5.6	0.46	8.3 ± 4.4	0.03	14.4 ± 7.1	-0.02	1.8 ± 1.9	0.83	2.6 ± 2.7	0.88	5.3 ± 7.6	0.94
	PBV (2014)	12.5 ± 9.6	0.09	9.1 ± 5.8	0.10	15.8 ± 7.2	0.08	3.6 ± 8.0	0.83	7.0 ± 7.9	0.65	9.4 ± 9.3	0.41
	2SR (2016)	7.5 ± 11.9	0.31	8.8 ± 6.4	0.03	12.9 ± 7.0	0.10	3.5 ± 7.8	0.85	3.6 ± 3.8	0.83	8.6 ± 9.8	0.75
	LAB (2016)	6.0 ± 6.6	0.37	8.5 ± 4.6	-0.03	16.2 ± 8.3	-0.01	2.6 ± 6.5	0.87	3.0 ± 2.6	0.73	7.5 ± 8.8	0.39
	POS (2017)	5.1 ± 6.2	0.58	7.8 ± 4.6	0.04	14.8 ± 7.3	0.04	2.2 ± 9.1	0.91	2.8 ± 2.7	0.89	3.7 ± 7.6	0.96
	LGI (2018)	4.7 ± 5.7	0.44	8.5 ± 3.8	0.02	14.2 ± 7.2	-0.09	4.2 ± 8.2	0.78	3.0 ± 2.4	0.72	4.8 ± 7.0	0.90
	OMIT (2022)	4.4 ± 5.7	0.53	7.5 ± 3.9	0.03	15.3 ± 5.9	-0.16	2.3 ± 9.1	0.88	3.3 ± 4.0	0.93	4.3 ± 6.1	0.95
Normalized (+ PPG Method)	GREEN (2008)	13.2 ± 9.7	0.23	13.1 ± 7.5	-0.03	18.7 ± 9.5	-0.16	12.4 ± 13.5	0.22	7.5 ± 4.4	0.46	9.0 ± 8.6	0.29
	ICA (2010)	10.9 ± 9.2	0.33	12.3 ± 5.1	-0.01	19.5 ± 6.2	0.12	3.5 ± 5.7	0.73	2.1 ± 0.5	0.66	3.2 ± 5.8	0.94
	PCA (2011)	11.9 ± 9.5	0.28	11.4 ± 5.1	0.02	19.4 ± 7.3	0.06	2.9 ± 2.9	0.53	2.6 ± 1.4	0.65	2.8 ± 2.9	0.90
	CHROM (2013)	6.9 ± 5.8	0.47	9.5 ± 4.3	-0.03	17.7 ± 9.3	-0.06	1.7 ± 2.1	0.78	1.5 ± 0.5	0.88	2.0 ± 2.8	0.95
	PBV (2014)	11.1 ± 8.3	0.11	12.1 ± 5.9	0.01	18.3 ± 10.1	0.03	10.3 ± 12.4	0.37	1.7 ± 0.7	0.88	4.3 ± 4.4	0.85
	2SR (2016)	12.0 ± 9.6	0.26	9.1 ± 4.6	0.03	15.2 ± 7.9	-0.03	1.3 ± 1.7	0.87	1.7 ± 0.6	0.95	1.7 ± 1.2	0.93
	LAB (2016)	8.0 ± 6.9	0.44	10.1 ± 4.7	-0.03	17.8 ± 11.3	-0.01	5.9 ± 7.5	0.49	1.9 ± 1.4	0.79	5.3 ± 5.7	0.62
	POS (2017)	8.7 ± 8.4	0.50	9.4 ± 4.8	0.10	18.1 ± 9.1	-0.07	1.5 ± 2.7	0.85	1.2 ± 0.4	0.94	1.4 ± 1.5	0.95
	LGI (2018)	8.8 ± 7.3	0.42	9.2 ± 4.3	-0.01	15.2 ± 7.0	0.01	2.6 ± 3.5	0.63	1.8 ± 0.6	0.76	3.6 ± 5.3	0.85
	OMIT (2022)	9.0 ± 8.0	0.39	10.6 ± 5.2	-0.06	14.2 ± 6.3	-0.08	2.2 ± 3.2	0.80	1.4 ± 0.5	0.91	3.4 ± 5.3	0.87
Multi-region (+ PPG Method)	GREEN (2008)	8.3 ± 7.3	0.26	9.1 ± 6.1	0.05	12.9 ± 11.0	0.21	5.4 ± 8.8	0.47	3.8 ± 2.4	0.59	1.4 ± 1.7	0.94
	ICA (2010)	9.2 ± 8.1	0.16	8.6 ± 4.3	0.03	14.3 ± 8.3	0.09	2.4 ± 2.6	0.71	1.5 ± 0.6	0.90	3.2 ± 3.3	0.75
	PCA (2011)	9.5 ± 9.7	0.22	8.6 ± 4.0	0.04	16.3 ± 8.1	0.17	2.4 ± 3.1	0.72	2.0 ± 1.0	0.75	3.5 ± 4.0	0.67
	CHROM (2013)	3.9 ± 3.6	0.58	8.8 ± 4.5	0.03	12.6 ± 6.2	0.10	1.2 ± 0.9	0.84	0.8 ± 0.5	0.98	1.5 ± 1.5	0.95
	PBV (2014)	9.5 ± 7.6	0.24	9.5 ± 4.6	0.09	15.9 ± 7.7	0.06	4.9 ± 7.4	0.53	3.7 ± 2.9	0.58	4.7 ± 5.2	0.65
	2SR (2016)	14.8 ± 8.6	0.08	9.1 ± 4.5	0.09	18.0 ± 8.7	0.06	5.7 ± 7.7	0.49	1.3 ± 0.5	0.91	6.6 ± 7.7	0.46
	LAB (2016)	6.4 ± 5.9	0.35	9.5 ± 4.4	-0.03	14.6 ± 7.1	0.03	3.0 ± 4.6	0.68	1.0 ± 0.5	0.96	2.5 ± 2.4	0.81
	POS (2017)	4.5 ± 3.3	0.57	8.0 ± 4.4	0.06	13.6 ± 6.7	0.05	1.4 ± 1.6	0.86	0.9 ± 0.4	0.96	0.9 ± 0.9	0.98
	LGI (2018)	4.5 ± 3.1	0.56	7.5 ± 3.5	0.11	10.8 ± 6.0	0.07	1.7 ± 1.9	0.72	1.4 ± 0.9	0.92	1.5 ± 1.6	0.96
	OMIT (2022)	4.4 ± 3.1	0.62	8.0 ± 4.1	0.08	9.3 ± 4.1	0.08	1.7 ± 2.6	0.86	0.8 ± 0.4	0.97	1.1 ± 1.2	0.98

The results for MAHNOB and UBFC2 are computed in a smaller subset of videos, according to [36]. The UBFC2 subset is comprised by 26 videos (Numbers 1, 3, 10, 11, 12, 13, 14, 15, 16, 17, 18, 20, 22, 23, 24, 25, 26, 27, 30, 31, 32, 33, 34, 35, 36 and 37). The MAHNOB subset is composed by 36 videos (Numbers 10, 1198, 1576, 2094, 2226, 2604, 2634, 2888, 3388, 3644, 556, 808, 1042, 1322, 1590, 2114, 2246, 2606, 274, 3006, 3396, 3662, 800, 814, 1184, 152, 1698, 2120, 2346, 2624, 276, 3136, 34, 408, 806 and 926). For LGI-PPGI, COHFACE, UBFC1 and PURE, we use the full set of videos. The bold values represent the “best value”. The best results achieved by each pipeline in their respective databases.

TABLE II
PERFORMANCE OF THE BASELINE AND MULTI-REGION PIPELINES (LGI-PPGI DATASET) FOR DIFFERENT HUMAN ACTIVITIES

Scenario	Baseline			Multi-region		
	MAE ± SD	PCC	RMSE	MAE ± SD	PCC	RMSE
Resting	1.8 ± 1.9	0.83	2.8	1.3 ± 1.3	0.87	1.7
Rotation	5.2 ± 3.2	0.10	9.4	5.1 ± 2.8	0.58	7.6
Talking	9.5 ± 9.8	0.59	11.3	5.8 ± 3.9	0.57	7.5
Gym	23.2 ± 11.4	0.16	27.6	7.5 ± 4.8	0.76	11.1

Coefficient improves, particularly in the gym scenario. Detailed results are presented in Table II.

C. Qualitative Results

We provide a visual representation of heart rate estimations from various pipelines and methods in Fig. 9. It shows the performance of four rPPG methods (CHROM, OMIT, POS, and GREEN) using the Multi-region pipeline on a PURE dataset video. CHROM, OMIT, and POS have similar performance, while GREEN struggles to track pulse rate

TABLE III
IMPACT OF THE NUMBER OF REGIONS IN rPPG EXTRACTION USING FACE2PPG-MULTI PIPELINE AND CHROM

Regions	LGI-PPGI			COHFACE		
	MAE ± SD	PCC	RMSE	MAE ± SD	PCC	RMSE
6 x 6	4.3 ± 4.3	0.60	5.9	8.9 ± 4.6	-0.03	11.0
7 x 7	4.1 ± 3.5	0.62	5.6	8.9 ± 4.4	0.01	11.0
8 x 8	4.0 ± 3.4	0.64	5.7	9.0 ± 4.3	0.03	11.1
9 x 9	3.9 ± 3.6	0.59	5.5	8.8 ± 4.5	0.03	11.0
10 x 10	4.0 ± 3.4	0.63	5.7	9.2 ± 4.4	-0.01	11.3
11 x 11	4.2 ± 4.0	0.61	5.8	9.3 ± 4.5	0.02	11.5
Forehead	5.4 ± 6.3	0.63	7.5	11.6 ± 5.5	-0.01	14.3
Cheeks	6.5 ± 5.5	0.42	8.7	10.9 ± 4.0	0.05	13.4
Combined	5.1 ± 4.5	0.60	7.3	10.1 ± 4.5	-0.01	12.6

during certain segments. The PCC metric reflects the similarity between the estimated (red) and ground-truth (blue) envelopes.

TABLE IV
COMPARISON OF *FACE2PPG-MULTI* PIPELINE WITH STATE OF THE ART SUPERVISED LEARNING-BASED (ORANGE) AND UNSUPERVISED NON-LEARNING-BASED (BLUE) METHODS

Metric	Method	Databases					
		LGI-PPGI	COHFACE	MAHNOB (*)	PURE	UBFC1	UBFC2 (*)
MAE (bpm \pm SD)	Face2PPG _{best} (Ours)	3.9 \pm 3.6	7.5 \pm 3.5	9.3 \pm 4.1	1.2 \pm 0.9	0.8 \pm 0.4	0.9 \pm 0.9
	Face2PPG _{omit} (Ours)	4.4 \pm 3.1	8.0 \pm 4.1	9.3 \pm 4.1	1.7 \pm 2.6	0.8 \pm 0.4	1.1 \pm 1.2
	CHROM (2013) [14]	9.4 \pm 12.6	12.4 \pm 7.0	18.6 \pm 9.0	1.6 \pm 2.0	2.2 \pm 0.8	8.3 \pm 10.9
	Li et al. (2014) [16]	-	19.9 \pm n/a	7.4 \pm n/a ¹	28.2 \pm n/a	-	-
	POS (2017) [19]	9.7 \pm 11.9	11.9 \pm 7.1	19.4 \pm 5.5	1.9 \pm 3.1	1.8 \pm 0.4	7.9 \pm 10.4
	FaceRPPG (2020) [21]	-	14.2 \pm 11.6	18.6 \pm 13.5	0.4 \pm 0.4 ²	-	2.7 \pm 6.7 ²
	HR-CNN (2018) [26]	-	8.1 \pm 9.5	7.3 \pm n/a	1.8 \pm n/a	-	4.9 \pm 5.1
	DeepPhys (2018) [27]	-	-	4.6 \pm n/a	0.8 \pm n/a	-	6.3 \pm 8.7
	rPPGNet (2019) [29]	-	-	5.5 \pm n/a	-	-	-
	Meta-rPPG (2020) [31]	-	9.3 \pm 11.5	3.0 \pm 4.9	-	-	6.0 \pm 7.1
	AutoHR (2020) [30]	-	-	3.8 \pm n/a	-	-	-
	PulseGAN (2021) [32]	-	-	4.2 \pm n/a	2.3 \pm n/a	-	2.1 \pm n/a
	AND-rPPG (2021) [33]	-	6.8 \pm 7.8	-	-	-	3.2 \pm 4.0
RMSE (bpm \pm SD)	Face2PPG _{best} (Ours)	5.5 \pm 5.0	9.8 \pm 4.6	12.5 \pm 5.8	1.8 \pm 1.9	1.1 \pm 0.5	2.8 \pm 3.0
	Face2PPG _{omit} (Ours)	6.4 \pm 4.9	10.5 \pm 4.4	12.5 \pm 5.8	2.9 \pm 4.6	1.1 \pm 0.5	3.1 \pm 3.8
	CHROM (2013) [14]	11.4 \pm 13.7	16.4 \pm 7.9	22.1 \pm 8.7	3.0 \pm 4.8	4.8 \pm 3.4	10.3 \pm 13.5
	Li et al. (2014) [16]	-	25.6 \pm n/a	7.6 \pm 6.9 ¹	31.0 \pm n/a	-	-
	POS (2017) [19]	11.9 \pm 13.0	16.1 \pm 8.1	27.7 \pm 6.2	3.6 \pm 6.3	2.6 \pm 0.9	9.9 \pm 11.7
	FaceRPPG (2020) [21]	-	-	-	-	-	-
	HR-CNN (2018) [26]	-	10.8 \pm n/a	9.2 \pm n/a	2.4 \pm n/a	-	5.9 \pm n/a
	DeepPhys (2018) [27]	-	-	-	1.5 \pm n/a	-	10.5 \pm n/a
	rPPGNet (2019) [29]	-	-	7.8 \pm 7.8	-	-	-
	Meta-rPPG (2020) [31]	-	12.3 \pm n/a	3.7 \pm n/a	-	-	7.4 \pm n/a
	AutoHR (2020) [30]	-	-	5.1 \pm 4.3	-	-	-
	PulseGAN (2021) [32]	-	-	6.5 \pm n/a	4.3 \pm n/a	-	4.4 \pm n/a
	AND-rPPG (2021) [33]	-	8.1 \pm n/a	-	-	-	4.8 \pm n/a

(*) We use the full set of each database for comparison, including MAHNOB (527 videos) and UBFC2 (42 videos). ¹ Li et al.'s [16] results were obtained using a smaller MAHNOB subset, removing low-quality PPG sections. Our fair configuration reproduction yields a MAE of 15.5. ² Gudi et al.'s [21] results use adaptive post-processing filtering. Their manuscript indicates an expected MAE of 1.2 for PURE and 4.1 for UBFC2 when using 10s. Some of the results for deep learning-based methods are referenced from [33] [70].

The bold values indicate again the “best value”. The best results attained for each individual database during the comparison of various rPPG methods.

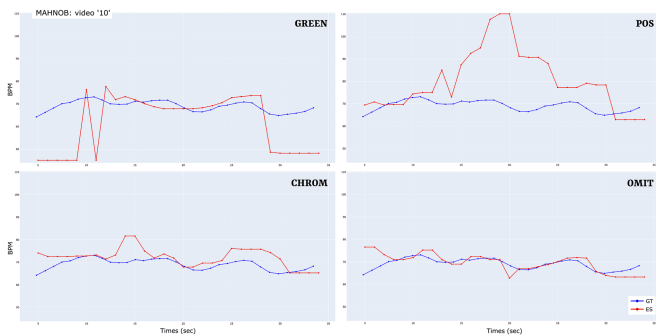


Fig. 10. Comparison of HR estimation using four rPPG methods with the Multi-region pipeline. Estimated heart rate (red) is extracted from the face, and the contact-based reference PPG signal (blue) is computed on a single MAHNOB video.

Fig. 10 shows the performance of four rPPG methods using the Multi-region pipeline on a MAHNOB database video. Despite the subject being static, high compression causes a loss of detail in raw RGB signals. CHROM and OMIT handle compression challenges better than GREEN and POS.

D. Evaluation of the Number of Regions

To evaluate the impact of the DMRS module in the Multi-region pipeline, we have designed a complementary experiment that measures how the results are affected depending on the number of facial regions used in the initial grid. We use CHROM as the RGB to rPPG conversion method, while all parameters remain the same except the number of initial available regions to select. In the comparison, in addition to region grids, we also include the typical fixed regions of the face, such as the forehead and cheeks.

For this experiment, the results are depicted in Table III. They show how generally, a moderately large number of regions results in smaller errors. The approach using fixed patches shows comparable results to configurations with low number of regions and proved to be still useful in some cases.

E. Comparison With the State of the Art

We compare our proposed pipeline with other supervised learning-based and unsupervised non learning-based methods as presented in the state of the art. We compare the results in terms of both MAE and RMSE and show them in Table IV. Table IV, shows two sets of results. Face2PPG-Multi_{best} results represent the best performance achieved using our multi-region pipeline

with various RGB to PPG transformation methods (CHROM, POS, LGI, or OMIT) showing the flexibility of our pipeline. Face2PPG-Multi_{omit} results display the outcomes applying just the OMIT method, to highlight its generalizability across diverse datasets.

It can be seen that our method, relying on the multi-region pipeline obtains better results than all unsupervised non-learning-based methods across all six benchmark datasets. Our results are also comparable to some recent supervised methods that require training on videos taken in similar conditions. Additionally, as shown in article [71], our method is highly efficient, taking only 17 ms per frame, and under 33 ms with face detection and alignment. This outperforms deep learning methods, which generally require longer processing times per frame, underscoring our approach's computational advantage.

VIII. CONCLUSION

In this article, we proposed a new supervised ("comparable results to other supervised methods") pipeline for the extraction of BVP signals from facial videos (rPPG). To enable a fair comparative evaluation among methods, we solved a set of smaller technical challenges such as problems with signal synchronization, use of different spectral analysis methods in extracted and reference signals, or inconsistent use of pipeline modules such as face detection and tracking or filtering. We proposed three novel contributions that improve the extraction of rPPG signals, especially in challenging conditions. First, we included a face normalization module, based on facial landmarks and a fixed triangle mesh that allowed the extraction of signals from exactly the same facial regions in a consistent manner. Second, we added the dynamic selection of facial regions that allowed to statistically discard those regions showing noise and artifacts. Finally, we proposed a novel RGB to PPG conversion method that increased the robustness of the extraction against compression artifacts. Our enhanced pipeline works in a purely unsupervised manner, and it is directly applicable in datasets collected in multiple conditions without any need of training data. The proposed pipeline achieves state-of-the-art results across multiple databases when compared with other unsupervised methods and shows comparable results to other supervised methods.

ACKNOWLEDGMENT

The authors appreciate the valuable feedback and suggestions provided by the associate editor, anonymous reviewers, and D.Sc. Matteo Pedone (CMVS - Oulu). Their insights and careful review greatly contributed to the clarity and accuracy of this article.

REFERENCES

- [1] A. B. Hertzman, "The blood supply of various skin areas as estimated by the photoelectric plethysmograph," *Amer. J. Physiol.-Legacy Content*, vol. 124, no. 2, pp. 328–340, 1938.
- [2] K. Shelley, S. Shelley, and C. Lake, "Pulse oximeter waveform: Photoelectric plethysmography," *Clin. Monit.: Practical Appl. Anesth. Crit. Care*, vol. 2, pp. 420–423, 2001.
- [3] J. Allen, "Photoplethysmography and its application in clinical physiological measurement," *Physiol. Meas.*, vol. 28, no. 3, pp. R1–39, Mar. 2007. [Online]. Available: <https://doi.org/10.1088/0967-3334/28/3/R01>
- [4] T. Tamura, "Current progress of photoplethysmography and spo2 for health monitoring," *Biomed. Eng. Lett.*, vol. 9, no. 2, pp. 21–36, 2019, doi: [10.1007/s13534-019-00097-w](https://doi.org/10.1007/s13534-019-00097-w).
- [5] P. Fonseca et al., "Validation of photoplethysmography-based sleep staging compared with polysomnography in healthy middle-aged adults," *Sleep*, vol. 40, no. 7, p. zsx097, 2017.
- [6] H. Kinnunen, A. Rantanen, T. Kentta, and H. Koskimaki, "Feasible assessment of recovery and cardiovascular health: Accuracy of nocturnal HR and HRV assessed via ring PPG in comparison to medical grade ECG," *Physiol. Meas.*, vol. 41, no. 4, p. 04NT01, 2020.
- [7] C. Alvarez Casado, P. Paananen, P. Siirtola, S. Pirttikangas, and M. Boddallo Lopez, "Meditation detection using sensors from wearable devices," in *Proc. Adjunct Proc. Int. Joint Conf. Pervasive Ubiquitous Comput.*, 2021, pp. 112–116.
- [8] D. R. Seshadri et al., "Wearable sensors for monitoring the internal and external workload of the athlete," *NPJ Digit. Med.*, vol. 2, 2019, Art. no. 71.
- [9] A. Ni, A. Azarang, and N. Kehtarnavaz, "A review of deep learning-based contactless heart rate measurement methods," *Sensors*, vol. 21, 2021, Art. no. 3719.
- [10] F. T.-Z. Khanam, A. Al-Naji, and J. Chahl, "Remote monitoring of vital signs in diverse non-clinical and clinical scenarios using computer vision systems: A review," *Appl. Sci.*, vol. 9, no. 20, 2019, Art. no. 4474.
- [11] W. Verkruysse, L. O. Svaasand, and J. S. Nelson, "Remote plethysmographic imaging using ambient light," *Opt. Exp.*, vol. 16, no. 26, pp. 21434–21445, Dec. 2008.
- [12] M.-Z. Poh, D. J. McDuff, and R. W. Picard, "Advancements in non-contact, multiparameter physiological measurements using a webcam," *IEEE Trans. Biomed. Eng.*, vol. 58, no. 1, pp. 7–11, Jan. 2011.
- [13] M. Lewandowska, J. Ruminski, T. Kocejko, and J. Nowak, "Measuring pulse rate with a webcam – a non-contact method for evaluating cardiac activity," in *Proc. IEEE Federated Conf. Comput. Sci. Inf. Syst.*, 2011, pp. 405–410.
- [14] G. de Haan and V. Jeanne, "Robust pulse rate from chrominance-based rPPG," *IEEE Trans. Biomed. Eng.*, vol. 60, no. 10, pp. 2878–2886, Oct. 2013.
- [15] G. Haan and A. Leest, "Improved motion robustness of remote-PPG by using the blood volume pulse signature," *Physiol. Meas.*, vol. 35, pp. 1913–1926, 2014.
- [16] X. Li, J. Chen, G. Zhao, and M. Pietikainen, "Remote heart rate measurement from face videos under realistic situations," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 4264–4271.
- [17] W. Wang, S. Stuijk, and G. de Haan, "A novel algorithm for remote photoplethysmography: Spatial subspace rotation," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 9, pp. 1974–1984, Sep. 2016.
- [18] Y. Yang, C. Liu, H. Yu, D. Shao, F. Tsow, and N. Tao, "Motion robust remote photoplethysmography in cielab color space," *J. Biomed. Opt.*, vol. 21, 2016, Art. no. 117001.
- [19] W. Wang, A. C. den Brinker, S. Stuijk, and G. de Haan, "Algorithmic principles of remote PPG," *IEEE Trans. Biomed. Eng.*, vol. 64, no. 7, pp. 1479–1491, Jul. 2017.
- [20] C. S. Pilz, S. Zaunseder, J. Krajewski, and V. Blazek, "Local group invariance for heart rate estimation from face videos in the wild," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops*, 2018, pp. 1335–1338.
- [21] A. Gudi, M. Bittner, and J. van Gemert, "Real-time webcam heart-rate and variability estimation with clean ground truth for evaluation," *Appl. Sci.*, vol. 10, no. 23, 2020, Art. no. 8630.
- [22] J. Gideon and S. Stent, "The way to my heart is through contrastive learning: Remote photoplethysmography from unlabelled video," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 3995–4004.
- [23] Z. Sun and X. Li, "Contrast-phys: Unsupervised video-based remote physiological measurement via spatiotemporal contrast," in *Proc. 17th Comput. Vis. Eur. Conf.*, 2022, pp. 492–510.
- [24] Y. Hsu, Y.-L. Lin, and W. Hsu, "Learning-based heart rate detection from remote photoplethysmography features," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2014, pp. 4433–4437.
- [25] X. Fan and J. Wang, "Bayesheart: A probabilistic approach for robust, low-latency heart rate monitoring on camera phones," in *Proc. 20th Int. Conf. Intell. User Interfaces*, 2015, pp. 405–416.
- [26] R. Špetlík, V. Franc, and J. Matas, "Visual heart rate estimation with convolutional neural network," in *Proc. Brit. Mach. Vis. Conf.*, 2018, pp. 3–6.
- [27] W. Chen and D. McDuff, "DeepPhys: Video-based physiological measurement using convolutional attention networks," in *Comput. Vis. ECCV 2018: 15th Eur. Conf.*, 2018, pp. 356–373, doi: [10.1007/978-3-030-01216-8_22](https://doi.org/10.1007/978-3-030-01216-8_22).

- [28] X. Niu, S. Shan, H. Han, and X. Chen, "Rhythmnet: End-to-end heart rate estimation from face via spatial-temporal representation," *IEEE Trans. Image Process.*, vol. 29, pp. 2409–2423, 2020.
- [29] Z. Yu, W. Peng, X. Li, X. Hong, and G. Zhao, "Remote heart rate measurement from highly compressed facial videos: An end-to-end deep learning solution with video enhancement," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 151–160.
- [30] Z. Yu, X. Li, X. Niu, J. Shi, and G. Zhao, "AutoHR: A strong end-to-end baseline for remote heart rate measurement with neural searching," *IEEE Signal Process. Lett.*, vol. 27, pp. 1245–1249, 2020.
- [31] E. Lee, E. Chen, and C.-Y. Lee, "Meta-rPPG: Remote heart rate estimation using a transductive meta-learner," in *Proc. 16th Comput. Vis. 16th Eur. Conf.*, 2020, pp. 392–409.
- [32] R. Song, H. Chen, J. Cheng, C. Li, Y. Liu, and X. Chen, "PulseGAN: Learning to generate realistic pulse waveforms in remote photoplethysmography," *IEEE J. Biomed. Health Inform.*, vol. 25, no. 5, pp. 1373–1384, May 2021.
- [33] B. Lokendra and G. Puneet, "AND-rPPG: A novel denoising-rPPG network for improving remote heart rate estimation," *Comput. Biol. Med.*, vol. 141, 2022, Art. no. 105146.
- [34] C.-H. Cheng, K.-L. Wong, J.-W. Chin, T.-T. Chan, and R. H. Y. So, "Deep learning methods for remote heart rate measurement: A review and future research agenda," *Sensors*, vol. 21, no. 18, 2021, Art. no. 6296.
- [35] Q. Zhan, W. Wang, and G. de Haan, "Analysis of CNN-based remote-PPG to understand limitations and sensitivities," *Biomed. Opt. Exp.*, vol. 11, no. 3, pp. 1268–1283, Mar. 2020.
- [36] G. Boccignone, D. Conte, V. Cuculo, A. D'Amelio, G. Grossi, and R. Lanzarotti, "An open framework for remote-PPG methods and their assessment," *IEEE Access*, vol. 8, pp. 216083–216103, 2020.
- [37] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Process. Lett.*, vol. 23, no. 10, pp. 1499–1503, Oct. 2016.
- [38] D. E. King, "DLIB-ML: A machine learning toolkit," *J. Mach. Learn. Res.*, vol. 10, pp. 1755–1758, 2009.
- [39] W. Liu et al., "SSD: Single shot multibox detector," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 21–37.
- [40] A. S. Sanchez-Moreno, J. Olivares-Mercado, A. Hernandez-Suarez, K. Toscano-Medina, G. Sanchez-Perez, and G. Benitez-Garcia, "Efficient face recognition system for operating in unconstrained environments," *J. Imag.*, vol. 7, no. 9, 2021, Art. no. 161.
- [41] V. Kazemi and J. Sullivan, "One millisecond face alignment with an ensemble of regression trees," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 1867–1874.
- [42] M. Kowalski, J. Naruniec, and T. Trzcinski, "Deep alignment network: A convolutional neural network for robust face alignment," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2017, pp. 2034–2043.
- [43] C. Álvarez Casado and M. Bordallo López, "Real-time face alignment: Evaluation methods, training strategies and implementation optimization," *J. Real-Time Image Process.*, vol. 18, no. 6, pp. 2239–2267, 2021.
- [44] H. Rakshit and M. A. Ullah, "A comparative study on window functions for designing efficient FIR filter," in *Proc. IEEE 9th Int. Forum Strategic Technol.*, 2014, pp. 91–96.
- [45] A. M. Unakafov, "Pulse rate estimation using imaging photoplethysmography: Generic framework and comparison of methods on a publicly available dataset," *Biomed. Phys. Eng. Exp.*, vol. 4, no. 4, Apr. 2018, Art. no. 045001.
- [46] B. Ly, E. Dyer, J. Feig, A. Chien, and S. Bino, "Research techniques made simple: Cutaneous colorimetry: A reliable technique for objective skin color measurement," *J. Invest. Dermatol.*, vol. 140, pp. 3–12, 2020.
- [47] X.-F. Teng and Y.-T. Zhang, "Theoretical study on the effect of sensor contact force on pulse transit time," *IEEE Trans. Biomed. Eng.*, vol. 54, no. 8, pp. 1490–1498, Aug. 2007.
- [48] H. Liu, J. Allen, S. G. Khalid, F. Chen, and D. Zheng, "Filtering-induced time shifts in photoplethysmography pulse features measured at different body sites: The importance of filter definition and standardization," *Physiol. Meas.*, vol. 42, no. 7, Jul. 2021, Art. no. 074001.
- [49] J. Fine et al., "Sources of inaccuracy in photoplethysmography for continuous cardiovascular monitoring," *Biosensors*, vol. 11, no. 4, 2021, Art. no. 126.
- [50] V. Hartmann, H. Liu, F. Chen, W. Hong, S. Hughes, and D. Zheng, "Toward accurate extraction of respiratory frequency from the photoplethysmogram: Effect of measurement site," *Front. Physiol.*, vol. 10, 2019, Art. no. 732.
- [51] A. Rodríguez and J. Ramos, "Video pulse rate variability analysis in stationary and motion conditions," *Biomed. Eng. Online*, vol. 17, 2018, Art. no. 11.
- [52] L. Nanni, A. Loreggia, A. Lumini, and A. Dorizza, "A standardized approach for skin detection: Analysis of the literature and case studies," *J. Imag.*, vol. 9, 2023, Art. no. 35.
- [53] A. Eke, P. Herman, L. Kocsis, and L. Kozák, "Fractal characterization of complexity in physiological temporal signals," *Physiol. Meas.*, vol. 23, pp. R31–R38, 2002.
- [54] C.-K. Peng, S. Havlin, H. Stanley, and A. Goldberger, "Quantification of scaling exponents and crossover phenomena in nonstationary heartbeat time series," *Chaos: An Interdiscipl. J. Nonlinear Sci.*, vol. 5, no. 1, pp. 82–87, 1995.
- [55] A. L. Goldberger et al., "Physiobank, physiotoolkit, and physionet," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000.
- [56] P. Stoica and R. Moses, *Spectral Analysis of Signals*. Hoboken, NJ, USA: Prentice Hall, 2005.
- [57] L. Sörnmo and P. Laguna, *Bioelectrical Signal Processing in Cardiac and Neurological Applications*. Amsterdam, The Netherlands: Elsevier Sci., 2005.
- [58] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Baltimore, MD, USA: JHU Press, 2013.
- [59] M. Shao, "Householder orthogonalization with a non-standard inner product," *SIAM J. Matrix Anal. Appl.*, vol. 44, no. 2, pp. 481–502, 2023, doi: [10.1137/21M1414814](https://doi.org/10.1137/21M1414814).
- [60] E. Anderson, Z. Bai, and J. Dongarra, "Generalized QR factorization and its applications," *Linear Algebra Appl.*, vol. 162–164, pp. 243–271, 1992.
- [61] Q. Liu, F. Davoine, J. Yang, Y. Cui, Z. Jin, and F. Han, "A fast and accurate matrix completion method based on QR decomposition and $L_{2,1}$ -norm minimization," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 30, no. 3, pp. 803–817, Mar. 2019.
- [62] M. Zimón et al., "An evaluation of noise reduction algorithms for particle-based fluid simulations in multi-scale applications," *J. Comput. Phys.*, vol. 325, pp. 380–394, 2016.
- [63] A. Sharma, K. Paliwal, S. Imoto, and S. Miyano, "Principal component analysis using QR decomposition," *Int. J. Mach. Learn. Cybern.*, vol. 4, no. 6, pp. 679–683, 2013.
- [64] R. Stricker, S. Müller, and H.-M. Gross, "Non-contact video-based pulse rate measurement on a mobile service robot," in *Proc. IEEE 23rd Int. Symp. Robot. Hum. Interactive Commun.*, 2014, pp. 1056–1062.
- [65] G. Heusch, A. Anjos, and S. Marcel, "A reproducible study on remote heart rate measurement," 2017, *arXiv:1709.00962*.
- [66] S. Bobbia, R. Macwan, Y. Benezeth, A. Mansouri, and J. Dubois, "Unsupervised skin tissue segmentation for remote photoplethysmography," *Pattern Recognit. Lett.*, vol. 124, pp. 82–90, 2019.
- [67] M. Soleymani, J. Lichtenauer, T. Pun, and M. Pantic, "A multimodal database for affect recognition and implicit tagging," *IEEE Trans. Affect. Comput.*, vol. 3, no. 1, pp. 42–55, Jan.–Mar. 2012.
- [68] J. Allen and A. Murray, "Effects of filtering on multisite photoplethysmography pulse waveform characteristics," in *Proc. IEEE Comput. Cardiol.*, 2004, pp. 485–488.
- [69] M. Rączynski, P. Werner, and A. Al-Hamadi, "Effects of video encoding on camera based heart rate estimation," *IEEE Trans. Biomed. Eng.*, vol. 66, no. 12, pp. 3360–3370, Dec. 2019.
- [70] A. K. Gupta, R. Kumar, L. Birla, and P. Gupta, "RADIANT: Better rPPG estimation using signal embeddings and transformer," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, 2023, pp. 4976–4986.
- [71] C. Álvarez Casado, M. L. Cañellas, and M. B. López, "Depression recognition using remote photoplethysmography from facial videos," *IEEE Trans. Affect. Comput.*, pp. 1–13, 2023, doi: [10.1109/TAFFC.2023.3238641](https://doi.org/10.1109/TAFFC.2023.3238641).