

Efficient Remote Photoplethysmography with Temporal Derivative Modules and Time-Shift Invariant Loss

Joaquim Comas¹, Adrià Ruiz², Federico Sukno¹

¹ Department of Information and Communication Technologies,
Pompeu Fabra University, Barcelona, Spain

² Seedtag, Madrid, Spain

¹ {joaquim.comas, federico.sukno}@upf.edu

² adriaruiz@seedtag.com

Abstract

We present a lightweight neural model for remote heart rate estimation focused on the efficient spatio-temporal learning of facial photoplethysmography (PPG) based on i) modelling of PPG dynamics by combinations of multiple convolutional derivatives, and ii) increased flexibility of the model to learn possible offsets between the facial video PPG and the ground truth. PPG dynamics are modelled by a Temporal Derivative Module (TDM) constructed by the incremental aggregation of multiple convolutional derivatives, emulating a Taylor series expansion up to the desired order. Robustness to ground truth offsets is handled by the introduction of TALOS (Temporal Adaptive Location Shift), a new temporal loss to train learning-based models. We verify the effectiveness of our model by reporting accuracy and efficiency metrics on the public PURE and UBFC-rPPG datasets. Compared to existing models, our approach shows competitive heart rate estimation accuracy with a much lower number of parameters and lower computational cost.

1. Introduction

Camera-based measurement of human physiological signals has gained considerable interest in the research community in recent years. The monitoring of vital signals such as heart rate (HR), heart rate variability (HRV), respiration rate (RR), oxygen saturation (SpO₂) or blood volume pulse (BVP) is crucial in the assessment of the physical and mental human state, and its potential applications [1, 13, 28]. Traditionally, physiological signals are acquired using continuous bio-signal monitoring devices, which typically require physical contact between the sensor and the skin. However, such physical contact limits the mobility of the subjects and may also be intrusive to users in many sce-

narios, causing some biases during the acquisition stage. Different approaches have emerged to estimate vital signals without physical contact to minimize these limitations, among which, video-based methods, especially remote photoplethysmography (rPPG), have attracted much attention due to their low cost, non-invasive nature, and wide applicability [26].

The rPPG or imaging photoplethysmography (iPPG) is the non-contact measurement of the PPG signal based on videos. The first rPPG measurements were proposed by Takano et al. [33] and Verkruysse et al. [36] in 2007 and 2008, respectively, who demonstrated the feasibility of extracting the pulse signal using a conventional camera. Since then, several handcrafted methods have been proposed [5, 37, 38] mainly based on optical models of the skin. These approaches are usually sensitive to non-controlled scenarios, including head motion and illumination changes, and require a multi-stage process where some of the steps are difficult to adjust.

The fast growth of deep learning (DL) techniques has led to more accurate and robust methods that have shown impressive results outperforming traditional approaches [25, 34, 42]. A key aspect for such success relates to the enhanced ability of these approaches to adequately model the spatial and, more importantly, temporal information (dynamics) present in the facial videos. Indeed, the most successful approaches to date are based on the use of 3D Convolutional Neural Networks (3DCNNs) that offer great flexibility to model any interaction between spatial and temporal patterns at once, though at the expense of considerably increasing the number of trainable parameters and thus the complexity and size of the model.

Alternatives to 3DCNNs have also been explored by combining spatial blocks based on 2DCNNs with additional modules or pre-processing operations that also take into account the temporal information of the signal. Notable ex-

amples include the use of normalized frame differences at the system’s input [3] or after a number of convolutional layers [40], as well as Temporal Shift Modules (TSM) [15] or combinations of 2D and 1D convolutions [14].

Even though the above alternatives avoid the use of the expensive 3D convolutions, achieving competitive performance in this way has proved challenging, requiring the clever combination of multiple modules to compensate for the noisy and unstable nature of rPPG dynamics, and resulting again in arguably large models. In this paper, we show that it is possible to model rPPG dynamics by means of local derivative filters, provided that they are applied to the appropriate spatial scale and to a sufficient order. This leads to a comparatively simpler architecture in which rPPG dynamics are captured by a chain of temporal derivative modules equipped with skip connections that emulates a Taylor series expansion up to a pre-established order. As shown in our experiments, our system is comparable to state-of-the-art approaches in terms of performance, but it is computationally more efficient and requires one or even two orders of magnitude less parameters.

In order to effectively exploit temporal information, another important aspect in deep learning approaches is the reliability of the ground-truth signals. While traditional techniques do not depend on the acquisition of the ground truth PPG because they directly estimate the rPPG signal from the skin pixels, data-driven methods need the ground-truth signal (e.g. from a pulsimeter) as a reference and assume that it is well synchronized with respect to the facial video. Nevertheless, these two modalities are often not correctly aligned, and there is an offset between the facial pulse signal and the ground truth signal, commonly caused by the Pulse Transit Time (PTT), whose exact value is unknown. A few recent works have highlighted the importance of such offset in performance, either by demonstrating improved results when the offset can be estimated beforehand and corrected [44] or by training models with modified loss functions that are more robust to the amount of ground-truth offset, although they do not model or estimate its value [6]. In contrast, we propose a more intuitive temporal-invariant loss by modifying the standard mean square error (MSE) loss to account for the temporal offset of the ground truth and automatically estimate its value for each training sequence.

1.1. Contribution

In this work, we present a lightweight architecture to address both the efficient modelling of rPPG dynamics and the impact of temporal desynchronization between video and ground-truth physiological data. Firstly, we introduce a novel architecture in which rPPG dynamics are captured by a chain of temporal derivative modules equipped with skip connections that emulates a Taylor series expansion up to

a pre-established order. Secondly, we present a new temporal loss, which we denote TALOS (Temporal Adaptive Location Shift), that allows training of deep learning methods invariantly to temporal offsets of the ground-truth signal. Compared to existing models, our approach shows competitive heart rate estimation accuracy with a much lower number of parameters and lower computational cost.

2. Related work

2.1. Camera-based Physiological measurement

Since Takano et al. [33] and Verkruysse et al. [36] evaluated the possibility of measuring HR remotely from facial videos, many researchers have proposed different methods to recover physiological data. Among them, some works consider regions of interest using various techniques, including Blind Source Separation [10, 26, 27], normalized Least Mean Squares [11] or self-adaptive matrix completion [35]. In contrast, other works rely on the skin optical reflection model by projecting all RGB skin pixels channels into a more refined subspace mitigating motion artifacts [5, 37, 38].

Recently, deep learning-based methods [9, 17, 23, 25, 31, 42] have outperformed conventional methods and achieved state-of-the-art performance estimating vital signs from facial videos. Some of these methods leverage prior knowledge learned from traditional methods and combine it with CNNs to exploit more sophisticated features [20, 22, 29]. On the other hand, some other researchers have aimed at fully end-to-end approaches [3, 25, 41]. Unlike previous methods, end-to-end models utilize facial videos as input to directly predict the rPPG signal.

2.2. Deep Spatio-Temporal modelling

Spatio-temporal modelling plays a crucial role in rPPG measurement. To learn better PPG features from facial video, DL approaches have been explored under different space-time schemes: sequential (e.g. CNN combined with LSTM [8, 9]), in parallel (e.g. two-branch CNN [3, 39]) or simultaneous (e.g. 3DCNNs [25, 41]). From these schemes, the most explored one is based on 3DCNNs, which simultaneously capture spatial and temporal information, reaching impressive results in heart rate (HR) estimation. However, they require high computational cost and a significant increase of trainable parameters compared with conventional 2DCNNs. Some alternatives appeared to overcome the computational cost of 3DCNNs while preserving the robustness of Spatio-temporal modelling. Firstly, Chen et al. [3] proposed a convolutional attention network (CAN) using normalized frame differences as input to emulate the first-order derivative in the temporal domain. Other studies attempted to boost 3DCNN temporal modelling by incorporating Temporal Shift Modules (TSM) [12, 15] and approx-

imating temporal representation employing a combination of 2D and 1D convolutions [14]. In contrast, Liu et al. [40] designed a novel Temporal Difference Convolution (TDC) to simulate the normalized frame difference from [3], but at feature level. However, we argue that this approach is not benefiting from higher orders dynamics that can contribute to more refined features from PPG signal recovery. A recent study proposed by Hill et al. [7] demonstrated the benefits of incorporating multi-derivative convolutions to obtain better PPG signal estimation. The problem of this method is that it is applying first and second-order derivatives as an input following the architecture presented by [3]. Unlike these approaches, we propose exploring the contribution of higher-order dynamics into feature level that is currently still largely unexplored.

2.3. Modelling temporal offsets in Deep Learning based approaches

One of the major drawbacks of Deep Learning (DL) approaches for rPPG compared to traditional ones such as [5, 37] is the need of a reliable ground truth of the BVP signal during the training process. More precisely, the rPPG estimation from facial videos demands a very precise synchronization between facial and physiological recordings. Nevertheless, since most of the existing databases employed finger pulse oximeter to record PPG signals [2, 21, 32, 46], they did not consider the physiological offset or PTT between the facial and finger blood flow [19]. Another potential synchronization error that can occur is during the acquisition procedure due to the rolling shutter effect. Mironenko et al. [18] proved the existence of a slight phase shift between both modalities caused by the progressive scanning of some imaging devices. Consequently, we are forcing DL models to learn inaccurate correspondences between the spatial and physiological information during the training stage.

Zhan et al. [44] were the first to determine the impact of phase-shift regarding the HR error between default and phase-corrected reference training labels. Since traditional approaches depend directly on the facial skin pixels, they used a POS approach to correct reference labels through the Hilbert Transform. Despite recent studies [30] are adopting the same strategy with different conventional approaches, this rectification implies an extra preprocessing step relying on the accurate PPG estimation of a traditional method. A different alternative to address the reference offset was presented by Gideon et al. [6], who introduced a frequency loss based on maximum cross-correlation (MCC) to be more robust to temporal offsets. However, they apply the cross-correlation loss at sub-sequence level, which provides too much flexibility to the model by ignoring the fact that temporal offsets are only dependent on the video or subject. In our ablation studies Sec. 4.3.2, we demonstrate that the

benefits of considering temporal offsets at the subject level instead of using a sub-sequence level also imply a better HR estimation.

In this paper, to overcome the phase-shift problem we propose an alternative and simpler loss based directly on the temporal domain, which provides very competitive results compared with the common losses used in rPPG estimation.

3. Methodology

In this section, we firstly introduce the details of our proposed rPPG end-to-end model. Subsequently, we focus on the presented objective function to deal the temporal offset between facial and physiological data.

3.1. Overview

The motivation of our proposed spatio-temporal model is three-fold: 1) projecting the input RGB images into an optimal latent space encoding the information of the rPPG signal while preserving the model’s computational complexity, 2) efficiently modelling temporal information using a cascaded temporal derivative module (TDM), to achieve the equivalent to a high-order Taylor approximation, 3) reduce the influence of PTT in the pulse estimation through T-LOS loss. Given an input facial RGB video sequence X with size $3 \times T \times H \times W$, the overall architecture is shown in Fig. 1. Here, 3 refers to RGB channels and T, H, W are the number of frames, height and width of each frame, respectively.

3.2. Spatio-temporal Modelling

The presented network is divided into the spatial encoder and the temporal derivative module. The encoder aims to learn relevant spatial filters that project the facial data into a space relevant to extracting pulsatile information. The encoder consists of two 2D convolutions with a 3x3 kernel size. Each convolution is followed by the hyperbolic tangent activation function, batch normalization (BN) and average pooling operation. Then, the rough spatial features of each frame are reshaped and zero-padded before being fed to the TDM to aggregate the information of all video frames. In our preliminary experiments, we observed that adding more spatial convolutional layers leads to a significantly higher computational complexity during inference without providing a relevant performance improvement. Our proposed TDM is composed of differential temporal convolutions (DTC) that model the dynamics of the spatial features by a chain of derivative filters up to the selected order, with the aim to approximate local dynamics as in a Taylor series expansion. Similar to [7], our architecture employs two consecutive DTC to model first and second temporal order derivatives which represent velocity and acceleration dynamics of rPPG. In order to show the benefits of modelling second-order temporal derivatives,

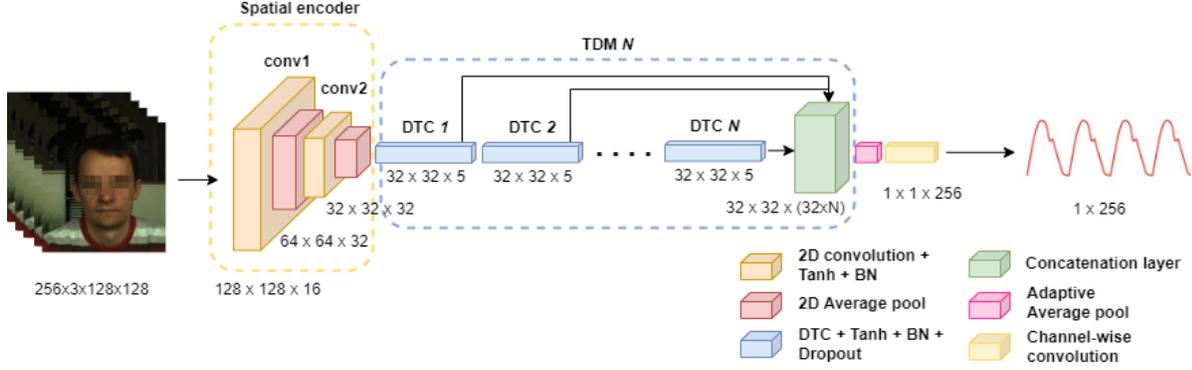


Figure 1. Overall structure of our proposed model.

we perform an ablation study where our model is able to extract more sophisticated features using second-order than zero and first order, which are the most common orders in prior work. To ensure this behavior, we implement the DTC blocks as 1D-convolutions with fixed weights i.e. non-trainable weights chosen to approximate a first-order derivative operator. Therefore, the first order DTC can be formulated as follows:

$$\text{DTC}_1(t) = \sum_{\tau \in R} w_{\partial t}(\tau) \cdot x(t + \tau) \quad (1)$$

where t denotes the current temporal location, x is the output tensor from the 2DCNN corresponding to the extracted spatial features, $w_{\partial t} = [-2, -1, 0, 1, 2]$ is the derivative filter using 5×1 kernel size and R is the receptive field. Our preliminary experiments revealed that the use of longer temporal windows did not lead to a significant performance improvement. Then, the n -th order DTC is defined recursively as:

$$\text{DTC}_n(t) = \sum_{\tau \in R} w_{\partial t}(\tau) \cdot \text{DTC}_{n-1}(t + \tau), \quad \forall n > 1 \quad (2)$$

where the temporal information is modeled using the output tensor of previous DTC instead of the spatial features from the 2D encoder. Later, we can define the TDM as the concatenation across the channel dimension of the tensor of all the DTC derivatives :

$$\text{TDM} = \text{Concat}(\text{DTC}_1, \text{DTC}_2, \dots, \text{DTC}_n) \quad (3)$$

Finally, the TDM output is projected into signal space using a channel-wise convolution operation with 1×1 kernel to generate the predicted rPPG signal of size T .

3.3. Loss function

The impact of PTT during the training process of learning-based methods has been previously discussed in

[44]. In particular, it has been shown that it is crucial to design an appropriate objective function robust to the temporal offset produced during the acquisition procedure or caused by the PTT between facial and finger regions. For this reason, we introduce the TALOS (Temporal Adaptive Location Shift) Loss, a new temporal loss to train deep learning models in the context of rPPG signal estimation.

To design our loss function, we consider two assumptions. Firstly, the time offset between facial data and the finger-PPG signal is in the order of milliseconds, and thus it is well bounded within $\pm \frac{F_s}{2}$ frames, where F_s is the video sampling frequency. Secondly, the PPG phase shift is the same for all video instances of the same subject because the time delay depends on each individual's physiology and the recording setting, which remains the same.

Considering the previous assumptions, we propose a loss function designed to be robust against possible unknown temporal offsets. In particular, we introduce a latent variable for each subject s as $K_s = \{k \mid \frac{k}{2} \in \mathbb{Z} : -\frac{F_s}{2} \leq k \leq \frac{F_s}{2}\}$, representing all the possible temporal shifts for a given subject. Given the previous definition, our goal is to learn a parametric distribution:

$$p_{\theta}(k|s) = \frac{\exp(\theta_k^s)}{\sum_{i=1}^K \exp(\theta_i^s)}, \quad (4)$$

where parameters $\theta^s \in \mathbb{R}^K$ encode the logits of a multinomial distribution for all the possible temporal offsets k given a subject s . Using the previously defined distributions, our proposed TALOS loss function is computed during optimization as described below.

Let $\hat{y}(t)$ be the predicted signal from our model, and $y(t)$ the ground-truth signal, both with the same length T . We zero-pad the ground truth signal with respect to all possible offsets k :

$$y_{pad} = \begin{cases} pad(k, y), & \text{if } k \leq 0 \\ pad(y, k), & \text{if } k > 0 \end{cases} \quad (5)$$

Then, the padded ground-truth signal y_{pad} is shifted for all possible delays within K and each y_k is used to compute the Mean Squared Error (MSE) with respect to the predicted signal y :

$$y_k = y_{pad}(t - k) \quad \forall k \in K \quad (6)$$

$$MSE(\hat{y}, y_k) = \frac{1}{T} \sum_{t=1}^T (\hat{y}(t) - y_k(t))^2 \quad (7)$$

Finally, we express our proposed TALOS loss function as:

$$\mathcal{L}_{\text{TALOS}} = \sum_{k \in K} MSE(\hat{y}, y_k) \cdot p_{\theta}(k|s), \quad (8)$$

where the MSE for each possible offset k is weighted according to the learned temporal-shift probability $p_{\theta}(k|s)$. Note that this probability is dependent on the the subject s corresponding to the specific ground-truth video.

The rationale of the above formulation is that, when the value of k is approximately the offset between the ground truth and the pulsatile variation captured by the input video, then the MSE is expected to reach its minimum value. However, k is not known beforehand and thus, we learn a distribution over a latent variable representing all the possible temporal shifts. Note that the parameters θ^s for this distributions are optimized during training and, as a consequence, we expect to assign a higher probability to offsets k leading to a lower MSE error.

4. Experiments

In this section, we first introduce the two benchmark datasets used in the experiments and describe the implementation details. Then, we compare the impact of different modules explained in Section 3 on the test results. Finally, we present a relative comparison with the existing rPPG methods.

4.1. Datasets

We evaluated our approach on the following two RGB video datasets.

PURE [32]: The Pulse Rate Detection Dataset contains 60 videos from 10 subjects (eight male, two female) performing six different head motion tasks: steady, talking, slow translation, fast translation, small rotation, and medium rotation. The facial videos were recorded using an ECO274CVGE camera with a resolution of 640 x 480 pixels and 30 frames per second (FPS). Each video is about 1 minute long and stored in uncompressed PNG format. The gold-standard measures of BVP and SpO2 were collected using a finger pulse oximeter (pulox CMS50E) with a sampling rate of 60 Hz. PURE contains predefined splits

for training and test (7 subjects for training, 3 subjects for testing), which we use to compare with the current related work.

UBFC-RPPG [2]: The UBFC-RPPG dataset includes 42 RGB videos from 42 subjects. The subjects were asked to play a time-sensitive mathematical game, emulating a standard human-computer interaction scenario, to obtain varied HR during the experiment. The recorded facial videos were acquired indoors with varying sunlight and indoor illumination at 30 FPS with a webcam (Logitech C920 HD Pro) at a resolution of 640x480 in uncompressed 8-bit RGB format. The bio-signals ground-truth were acquired using a CMS50E transmissive pulse oximeter to record PPG signal and heart rate with a 60 Hz sampling rate. To validate UBFC, we use the first 28 subjects for training and the last 14 subjects for testing following [16].

4.2. Implementation details

4.2.1 Preprocessing and training procedure

We adopt the same preprocessing stage for each dataset in all our experiments. Firstly, we estimate facial regions for each frame using the MTCNN [45] algorithm adding a 50% scale size of the detected box and resizing each frame to 128×128 pixels. The ground-truth bio-signal is preprocessed following [4] to denoise the raw PPG signal, which facilitates a better model convergence during the training procedure. The resulting bio-signal is then downsampled to 30 Hz so that it matches the sampling rate of the videos.

We implemented our model using Pytorch 1.8.0 [24] and trained it on a single NVIDIA GTX1080Ti. We utilized sequences of 256 frames without overlap and Adadelta optimizer during training [43]. In addition, when using our TALOS loss, we incorporated an extra SGD optimizer with a 0.01 learning rate to optimize the parameters θ^s of the temporal-shifts distributions for each subject. Finally, the estimated HR was computed from the predicted rPPG signal using the power spectral density (PSD). Before calculating the HR value, we applied a band-pass filter with cutoff frequencies of 0.75Hz and 2.5 Hz, similarly to [3].

4.2.2 Metrics and evaluation

To evaluate the HR estimation performance of the proposed model, we adopted the same metrics used in the literature, such as the mean absolute HR error (MAE), the root mean squared HR error (RMSE) and Pearson’s correlation coefficients R [11], [35]. Besides, we computed the computational performance in terms of the number of parameters and Multiple-Accumulate (MACs) ¹ to compare the efficiency of our model with respect to the major existing

¹<https://github.com/sovrasov/flops-counter-pytorch>

methods. Ablation experiments were performed using 256-frame sequences (8.5 seconds approx) with no overlap to compute HR estimation and all reported metrics, which is more challenging and informative than HR estimation based on the whole video sequence at once. Nevertheless, we also computed whole-video performance to fairly compare our method to the state-of-the-art, because most prior work adopts whole-video evaluation. Thus, we compare our approach with different traditional and deep learning methods. Since not all learning-based methods report HR results on the same dataset, we re-implemented the PhysNet model [41] using the code provided by the author² and the hyper-parameters in the paper. For DeepPhys [3] and Zhan et al. [44] model, the source code is not publicly available, so we also re-implemented them according to the details provided in the respective papers.

4.3. Ablation studies

4.3.1 Evaluation of Spatio-temporal modelling

Our first experiment to configure our model analyses the influence of aggregating DTCs. For this reason, we evaluate our model into three different schemes. Firstly, we consider our model using only spatial information without modelling temporal dynamics. The 2DCNN spatial encoder forms this first configuration, also called TDM_0 , since it represents a zero-order derivative. Later, to create high-order models, we introduce sequential DTCs creating TDM_1 and TDM_2 , representing the first and second-order derivatives. Tab. 1 summarized the sequence HR estimation for each temporal derivative order, which indicates that the aggregation of high-order derivatives is producing a more robust spatio-temporal representation for rPPG.

Table 1. Evaluation of TDC aggregation

Derivative order	Sequence evaluation		
	MAE↓	RMSE↓	R↑
TDM_0	4.34	8.05	0.80
TDM_1	2.84	6.01	0.88
TDM_2	2.68	5.55	0.90

As mentioned in Sec. 2.2, there are different approaches to model spatio-temporal information. Then, in this second experiment, we fix the first layers of our model (spatial-only), used as a baseline model, and we evaluate different choices for the temporal modelling stage. To this end, we implement two modified versions of our proposed model. In the first version, we adopt the best configuration found in the first experiment, which consists of TDM_2 , shown in Fig. 1, while in the second version, we combine spatial and temporal information simultaneously, substituting

²<https://github.com/ZitongYu/PhysNet>

the TDM by two 3D convolutional layers. To validate the capability of the different configurations, we evaluate each architecture in terms of sequence HR estimation, but also the performance of each architecture, summarized in Tab. 2. In terms of HR estimation, our proposed TDM and the modified 3D architecture obtained similar results, while the 2DCNN model produces higher HR error than the other configurations. On the other hand, regarding the computational cost, we observe a significant difference in the number of parameters and MACs between 2DCNN and our TDM approach compared with the 3DCNN configuration. This suggests that constrained temporal modelling based on dynamic properties can capture temporal information as accurately as 3DCNNs but with far less parameters and reduced computational cost, similar to 2DCNN performance.

Table 2. Evaluation of spatio-temporal modelling.

Model	Heart Rate			Performance	
	MAE↓	RMSE↓	R↑	Params(K)	MACs(G)
2DCNN	4.34	8.05	0.80	5.17	6.95
3DCNN	2.69	5.66	0.89	15.51	9.70
TDM	2.68	5.55	0.90	5.26	7.08

4.3.2 Evaluation of the objective function

To evaluate the effectiveness of TALOS loss, we compare it to other loss functions under the same network configuration explained in the previous section. To this end, we choose the MSE, and Negative Pearson Correlation (NPC) losses [41], which are widely used in the prior work for remote heart rate estimation. Furthermore, to assess the robustness of temporal offsets on the ground-truth signal, we also evaluate a frequency domain loss function, the maximum cross-correlation (MCC) loss proposed in [6].

Table 3. Comparison of different loss functions.

Loss	Sequence evaluation		
	MAE↓	RMSE↓	R↑
MSE	2.68	5.55	0.90
NPC [41]	2.67	5.62	0.89
MCC [6]	2.44	3.73	0.96
TALOS	2.33	3.41	0.96

Tab. 3 summarizes the results, which show a similar trend between MSE and NPC losses. This can be explained because NPC is only focused on the similarity between signals without considering the morphology and range of the signals, which can lead to some estimation error. Usually, the majority of methods that adopt NPC apply some previous normalization before propagating the loss, which we do not consider to compare all the loss functions in a more

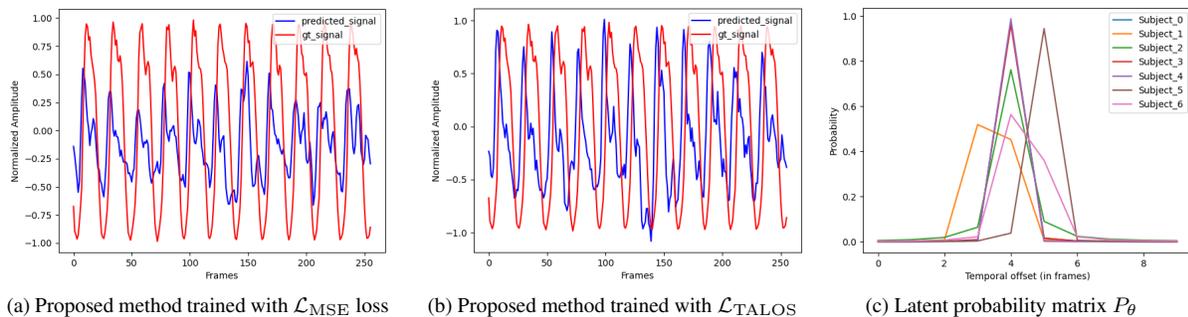


Figure 2. PPG sequence estimation on individual’s PURE test data using the standard MSE loss, Fig. 2a, and the proposed TALOS loss, Fig. 2b. In Fig. 2b we observe how the model is able to predict the PPG signal in the temporal instance where the regression error is minimum, regardless of the ground-truth phase. On Fig. 2c is shown the learned frame offset for each subject during the training process using TALOS loss. The offset range is limited between 0 and 300 ms for better visualization.

direct and fair manner. In contrast, we can appreciate a significant difference between these previous temporal loss functions and temporal offset invariant losses such as TALOS or MCC. These results show the impact on HR estimation of the PTT delay between the ground truth and the facial data, which makes it harder for the model to learn the actual relation between the predicted signal from the skin color changes and the label signal when using conventional loss functions. Finally, comparing MCC with TALOS, we observe that our proposed loss obtains better results in MAE and RMSE. This difference can be explained because MCC considers each sequence or video independently while TALOS loss enforces to learn a consistent shift distribution for each subject. In Fig. 2, we can note that the predicted signal amplitude, shown in Fig. 2b, is substantially more accurate than the estimates produced under standard losses (Fig. 2a), while the offset between the ground truth and the prediction is still clearly visible. This effect is due the built-in estimate of the shift probabilities of TALOS, which we show in Fig. 2c: notice how the estimated shift probabilities for each subject always follow a similar pattern dominated by a peak that spreads one or two frames that is surrounded by near-zero values elsewhere. Further, these peaks are localized between 3 to 5 frames, suggesting offsets that range from 100 to 170 milliseconds, depending on the subject.

4.4. Comparison with existing methods

After the preliminary studies from the previous sections, we select our best model configuration to compare against other reported results on the PURE and UBFC-rPPG datasets. Besides, we compare the efficiency of our model with respect to the reproduced methods explained on Sec. 4.2. In contrast to the ablation studies of Sec. 4.3, we estimate the average HR for each video as a whole, to allow for fair comparison to results from other researchers. Results are summarized in Tab. 4.

1) Results on PURE dataset: We first conduct the video

HR evaluation on the PURE dataset. These results show that DL methods like HR-CNN and DeepPhys perform much better than the hand-crafted methods such as POS and CHROM. This suggests that learning-based models can extract more representative features than hand-crafted methods for PPG signal and HR estimation. Indeed, the results of the compared DL approaches vary between 2.29 and 2.90 BPM RMSE, which denotes the good performance in non-compressed data and controlled scenarios such as the PURE dataset. Regarding our proposed method, we can appreciate that it outperforms some DL methods such as HR-CNN or PhysNet, achieving at the same time similar results as Zhan et al. [44] and DeepPhys, which obtains the best results.

2) Results on UBFC-rPPG dataset: We also evaluate our model in the public UBFC dataset. Comparing these results to those from the PURE dataset, we observe that UBFC is more challenging since the RMSE from the different methods varies between 3.08 and 6.61 BPM. It can be seen that our proposed model achieves the best HR video results concerning both traditional and DL approaches. This performance is possible through TALOS loss, which considerably improves our baseline results by adopting a standard MSE regression. Comparing the impact of TALOS loss in HR video estimation for PURE and UBFC datasets, we note a significant error decrease in UBFC and a more modest improvement in the PURE dataset. The limited number of subjects can explain this difference in the PURE dataset, which consists of several recordings from only ten subjects, while UBFC datasets contain more subject diversity and fewer recordings per user.

3) Results on computation cost: Finally, we analyze the performance efficiency of our method compared with the publicly available and re-implemented DL methods. Tab. 5 shows that our lightweight model is extremely efficient with respect to the state-of-the-art models in terms of parameters and MACs. Considering this reduced computational cost, our model outperforms the PhysNet, consisting of a

Table 4. : Performance of heart rate measurement for PURE and UBFC-rPPG datasets.

Method	PURE			UBFC		
	MAE↓	RMSE↓	R↑	MAE↓	RMSE↓	R↑
CHROM [5]	2.07	2.50	0.99	3.44	4.61	0.97
POS [37]	3.14	10.57	0.95	2.44	6.61	0.94
HR-CNN [31]	1.84	2.37	0.98	-	-	-
DeepPhys [3]	1.84	2.31	0.99	2.90	3.63	0.98
PhysNet [41]	2.16	2.70	0.99	2.95	3.67	0.98
Zhan et al. [44]	1.82	2.29	0.99	2.44	3.17	0.99
Gideon et al. [6]	2.30	2.90	0.99	-	-	-
AND-rPPG [16]	-	-	-	2.67	4.07	0.96
Ours	1.89	2.33	0.99	2.54	3.31	0.99
Ours + TALOS	1.83	2.30	0.99	2.32	3.08	0.99

3DCNN Spatio-temporal auto-encoder. This means that a reduced model with well-designed but constrained temporal modelling, and suitable PPG preprocessing can achieve similar or better results than 3DCNNs, which usually require massive data to train all the model parameters.

Table 5. Efficiency performance on tested methods.

Methods	Params(K)	MACs(G)
DeepPhys [3]	1060	16.44
PhysNet [41]	768	112.33
Zhan et al. [44]	2620	59.40
Ours	5.26	7.080

5. Conclusions

In this paper, we propose a lightweight model for remote heart rate measurement. The presented method effectively estimates rPPG spatio-temporal features by aggregating multiple temporal derivative filters up to the desired order. Moreover, to mitigate the effect of possible temporal offsets such as PTT, we introduce a new objective loss function, which we call TALOS, designed to learn the latent shift probabilities of the ground truth during the training process, which allows to minimize the error between the predicted and label signals under the estimated relative offset. Through experiments on the PURE and UBFC-rPPG datasets, our proposed framework demonstrates a competitive HR performance with reduced computational requirements, which facilitates applicability to real/low-resource scenarios. In our future work, we aim to explore the combination of higher-order dynamics to further improve the temporal modelling for rPPG estimation.

Acknowledgments

This work is partly supported by the eSCANFace project (PID2020-114083GB-I00) funded by the Spanish Ministry

of Science and Innovation.

References

- [1] Yannick Benezeth, Peixi Li, Richard Macwan, Keisuke Nakamura, Randy Gomez, and Fan Yang. Remote heart rate variability for emotional state monitoring. In *2018 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*, pages 153–156. IEEE, 2018. 1
- [2] Serge Bobbia, Richard Macwan, Yannick Benezeth, Alamin Mansouri, and Julien Dubois. Unsupervised skin tissue segmentation for remote photoplethysmography. *Pattern Recognition Letters*, 124:82–90, 2019. 3, 5
- [3] Weixuan Chen and Daniel McDuff. Deepphys: Video-based physiological measurement using convolutional attention networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 349–365, 2018. 2, 3, 5, 6, 8
- [4] Lorenzo Dall’Olio, Nico Curti, Daniel Remondini, Yosef Safi Harb, Folkert W Asselbergs, Gastone Castellani, and Hae-Won Uh. Prediction of vascular aging based on smartphone acquired ppg signals. *Scientific reports*, 10(1):1–10, 2020. 5
- [5] Gerard De Haan and Vincent Jeanne. Robust pulse rate from chrominance-based rppg. *IEEE Transactions on Biomedical Engineering*, 60(10):2878–2886, 2013. 1, 2, 3, 8
- [6] John Gideon and Simon Stent. The way to my heart is through contrastive learning: Remote photoplethysmography from unlabelled video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3995–4004, 2021. 2, 3, 6, 8
- [7] Brian L Hill, Xin Liu, and Daniel McDuff. Learning higher-order dynamics in video-based cardiac measurement. *arXiv preprint arXiv:2110.03690*, 2021. 3
- [8] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [9] Eugene Lee, Evan Chen, and Chen-Yi Lee. Meta-rppg: Remote heart rate estimation using a transductive meta-learner. In *European Conference on Computer Vision*, pages 392–409. Springer, 2020. 2

- [10] Magdalena Lewandowska, Jacek Rumiński, Tomasz Kocejko, and Jędrzej Nowak. Measuring pulse rate with a webcam—a non-contact method for evaluating cardiac activity. In *2011 federated conference on computer science and information systems (FedCSIS)*, pages 405–410. IEEE, 2011. [2](#)
- [11] Xiaobai Li, Jie Chen, Guoying Zhao, and Matti Pietikainen. Remote heart rate measurement from face videos under realistic situations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4264–4271, 2014. [2](#), [5](#)
- [12] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7083–7093, 2019. [2](#)
- [13] Siqi Liu, Pong C Yuen, Shengping Zhang, and Guoying Zhao. 3d mask face anti-spoofing with remote photoplethysmography. In *European Conference on Computer Vision*, pages 85–100. Springer, 2016. [1](#)
- [14] Si-Qi Liu and Pong C Yuen. A general remote photoplethysmography estimator with spatiotemporal convolutional network. In *15th IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, pages 685–692. IEEE Computer Society, 2020. [2](#), [3](#)
- [15] Xin Liu, Josh Fromm, Shwetak Patel, and Daniel McDuff. Multi-task temporal shift attention networks for on-device contactless vitals measurement. *Advances in Neural Information Processing Systems*, 33:19400–19411, 2020. [2](#)
- [16] Birla Lokendra and Gupta Puneet. And-rppg: A novel denoising-rppg network for improving remote heart rate estimation. *Computers in biology and medicine*, 141:105146, 2022. [5](#), [8](#)
- [17] Hao Lu, Hu Han, and S Kevin Zhou. Dual-gan: Joint bvp and noise modeling for remote physiological measurement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12404–12413, 2021. [2](#)
- [18] Yuriy Mironenko, Konstantin Kalinin, Mikhail Kopeliovich, and Mikhail Petrushan. Remote photoplethysmography: Rarely considered factors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 296–297, 2020. [3](#)
- [19] Andreia Vieira Moço, Sander Stuijk, and Gerard de Haan. Skin inhomogeneity as a source of error in remote ppg-imaging. *Biomedical optics express*, 7(11):4718–4733, 2016. [3](#)
- [20] Xuesong Niu, Hu Han, Shiguang Shan, and Xilin Chen. Synrhythm: Learning a deep heart rate estimator from general to specific. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 3580–3585. IEEE, 2018. [2](#)
- [21] Xuesong Niu, Hu Han, Shiguang Shan, and Xilin Chen. Vipl-hr: A multi-modal database for pulse estimation from less-constrained face video. In *Asian Conference on Computer Vision*, pages 562–576. Springer, 2018. [3](#)
- [22] Xuesong Niu, Shiguang Shan, Hu Han, and Xilin Chen. Rhythmnet: End-to-end heart rate estimation from face via spatial-temporal representation. *IEEE Transactions on Image Processing*, 29:2409–2423, 2019. [2](#)
- [23] Ewa M Nowara, Daniel McDuff, and Ashok Veeraraghavan. The benefit of distraction: Denoising camera-based physiological measurements using inverse attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4955–4964, 2021. [2](#)
- [24] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. [5](#)
- [25] Olga Perepelkina, Mikhail Artemyev, Marina Churikova, and Mikhail Grinenko. Hearttrack: Convolutional neural network for remote video-based heart rate monitoring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 288–289, 2020. [1](#), [2](#)
- [26] Ming-Zher Poh, Daniel J McDuff, and Rosalind W Picard. Advancements in noncontact, multiparameter physiological measurements using a webcam. *IEEE transactions on biomedical engineering*, 58(1):7–11, 2010. [1](#), [2](#)
- [27] Ming-Zher Poh, Daniel J McDuff, and Rosalind W Picard. Non-contact, automated cardiac pulse measurements using video imaging and blind source separation. *Optics express*, 18(10):10762–10774, 2010. [2](#)
- [28] Vincenzo Ronca, Andrea Giorgi, Dario Rossi, Antonello Di Florio, Gianluca Di Flumeri, Pietro Aricò, Nicolina Sciarraffa, Alessia Vozzi, Luca Tamborra, Ilaria Simonetti, et al. A video-based technique for heart rate and eye blinks rate estimation: A potential solution for telemonitoring and remote healthcare. *Sensors*, 21(5):1607, 2021. [1](#)
- [29] Rencheng Song, Huan Chen, Juan Cheng, Chang Li, Yu Liu, and Xun Chen. PulseGAN: Learning to generate realistic pulse waveforms in remote photoplethysmography. *IEEE Journal of Biomedical and Health Informatics*, 25(5):1373–1384, 2021. [2](#)
- [30] Jeremy Speth, Nathan Vance, Patrick Flynn, Kevin Bowyer, and Adam Czajka. Unifying frame rate and temporal dilations for improved remote pulse detection. *Computer Vision and Image Understanding*, 210:103246, 2021. [3](#)
- [31] Radim Špetlík, Vojtech Franc, and Jirí Matas. Visual heart rate estimation with convolutional neural network. In *Proceedings of British Machine Vision Conference*, 2018. [2](#), [8](#)
- [32] Ronny Stricker, Steffen Müller, and Horst-Michael Gross. Non-contact video-based pulse rate measurement on a mobile service robot. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pages 1056–1062. IEEE, 2014. [3](#), [5](#)
- [33] Chihiro Takano and Yuji Ohta. Heart rate measurement based on a time-lapse image. *Medical engineering & physics*, 29(8):853–857, 2007. [1](#), [2](#)
- [34] Yun-Yun Tsou, Yi-An Lee, Chiou-Ting Hsu, and Shang-Hung Chang. Siamese-rppg network: remote photoplethysmography signal estimation from face videos. In *Proceedings of the 35th Annual ACM Symposium on Applied Computing*, pages 2066–2073, 2020. [1](#)
- [35] Sergey Tulyakov, Xavier Alameda-Pineda, Elisa Ricci, Lijun Yin, Jeffrey F Cohn, and Nicu Sebe. Self-adaptive ma-

- trix completion for heart rate estimation from face videos under realistic conditions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2396–2404, 2016. 2, 5
- [36] Wim Verkruijse, Lars O Svaasand, and J Stuart Nelson. Remote plethysmographic imaging using ambient light. *Optics express*, 16(26):21434–21445, 2008. 1, 2
- [37] Wenjin Wang, Albertus C den Brinker, Sander Stuijk, and Gerard De Haan. Algorithmic principles of remote ppg. *IEEE Transactions on Biomedical Engineering*, 64(7):1479–1491, 2016. 1, 2, 3, 8
- [38] Wenjin Wang, Sander Stuijk, and Gerard De Haan. A novel algorithm for remote photoplethysmography: Spatial sub-space rotation. *IEEE transactions on biomedical engineering*, 63(9):1974–1984, 2015. 1, 2
- [39] Zhi-Kuan Wang, Ying Kao, and Chiou-Ting Hsu. Vision-based heart rate estimation via a two-stream cnn. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 3327–3331. IEEE, 2019. 2
- [40] Zitong Yu, Xiaobai Li, Xuesong Niu, Jingang Shi, and Guoying Zhao. Autohr: A strong end-to-end baseline for remote heart rate measurement with neural searching. *IEEE Signal Processing Letters*, 27:1245–1249, 2020. 2, 3
- [41] Z. Yu, Xiao-Bai Li, and G. Zhao. Remote photoplethysmograph signal measurement from facial videos using spatio-temporal networks. In *BMVC*, 2019. 2, 6, 8
- [42] Zitong Yu, Wei Peng, Xiaobai Li, Xiaopeng Hong, and Guoying Zhao. Remote heart rate measurement from highly compressed facial videos: an end-to-end deep learning solution with video enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 151–160, 2019. 1, 2
- [43] Matthew D Zeiler. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012. 5
- [44] Qi Zhan, Wenjin Wang, and Gerard de Haan. Analysis of cnn-based remote-ppg to understand limitations and sensitivities. *Biomedical optics express*, 11(3):1268–1283, 2020. 2, 3, 4, 6, 7, 8
- [45] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE signal processing letters*, 23(10):1499–1503, 2016. 5
- [46] Zheng Zhang, Jeff M Girard, Yue Wu, Xing Zhang, Peng Liu, Umur Ciftci, Shaun Canavan, Michael Reale, Andy Horowitz, Huiyuan Yang, et al. Multimodal spontaneous emotion corpus for human behavior analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3438–3446, 2016. 3