



# PhysFormer++: Facial Video-Based Physiological Measurement with SlowFast Temporal Difference Transformer

Zitong Yu<sup>1</sup> · Yuming Shen<sup>2</sup> · Jingang Shi<sup>3</sup> · Hengshuang Zhao<sup>4</sup> · Yawen Cui<sup>5</sup> · Jiehua Zhang<sup>5</sup> · Philip Torr<sup>2</sup> · Guoying Zhao<sup>5</sup>

Received: 27 April 2022 / Accepted: 6 January 2023 / Published online: 15 February 2023  
© The Author(s) 2023

## Abstract

Remote photoplethysmography (rPPG), which aims at measuring heart activities and physiological signals from facial video without any contact, has great potential in many applications (e.g., remote healthcare and affective computing). Recent deep learning approaches focus on mining subtle rPPG clues using convolutional neural networks with limited spatio-temporal receptive fields, which neglect the long-range spatio-temporal perception and interaction for rPPG modeling. In this paper, we propose two end-to-end video transformer based architectures, namely PhysFormer and PhysFormer++, to adaptively aggregate both local and global spatio-temporal features for rPPG representation enhancement. As key modules in PhysFormer, the temporal difference transformers first enhance the quasi-periodic rPPG features with temporal difference guided global attention, and then refine the local spatio-temporal representation against interference. To better exploit the temporal contextual and periodic rPPG clues, we also extend the PhysFormer to the two-pathway SlowFast based PhysFormer++ with temporal difference periodic and cross-attention transformers. Furthermore, we propose the label distribution learning and a curriculum learning inspired dynamic constraint in frequency domain, which provide elaborate supervisions for PhysFormer and PhysFormer++ and alleviate overfitting. Comprehensive experiments are performed on four benchmark datasets to show our superior performance on both intra- and cross-dataset testings. Unlike most transformer networks needed pretraining from large-scale datasets, the proposed PhysFormer family can be easily trained from scratch on rPPG datasets, which makes it promising as a novel transformer baseline for the rPPG community.

**Keywords** RPPG · Temporal difference transformer · SlowFast · Cross-attention · Periodic-attention

---

Communicated by Jingdong Wang.

---

✉ Guoying Zhao  
guoying.zhao@oulu.fi

Zitong Yu  
zitong.yu@ieee.org

Yuming Shen  
ym\_zmxncbv@hotmail.com

Jingang Shi  
jingang@xjtu.edu.cn

Hengshuang Zhao  
hszhao@cs.hku.hk

Yawen Cui  
yawen.cui@oulu.fi

Jiehua Zhang  
jiehua.zhang@oulu.fi

Philip Torr  
philip.torr@eng.ox.ac.uk

## 1 Introduction

Physiological signals such as heart rate (HR), respiration frequency (RF), and heart rate variability (HRV) are important vital signs to be measured in many circumstances, especially for healthcare or medical purposes. Traditionally, the electrocardiography (ECG) and photoplethysmograph (PPG) or blood volume pulse (BVP) are the two most common ways for measuring heart activities and corresponding physiological signals. However, both ECG and PPG/BVP sensors need to be attached to body parts, which may cause discomfort

<sup>1</sup> Great Bay University, Dongguan 523000, China

<sup>2</sup> University of Oxford, Oxford OX13PJ, UK

<sup>3</sup> Xi'an Jiaotong University, Xi'an 710049, China

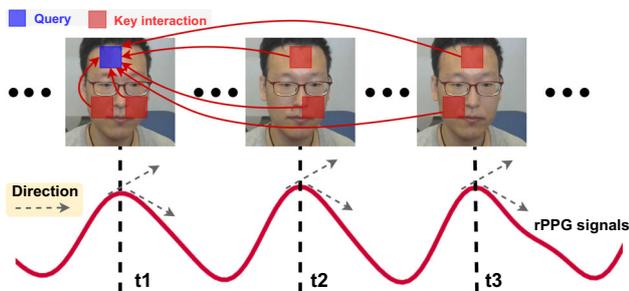
<sup>4</sup> The University of Hong Kong, Hong Kong, China

<sup>5</sup> University of Oulu, 90014 Oulu, Finland

and are inconvenient for long-term monitoring. To counter this issue, remote photoplethysmography (rPPG) (Chen et al., 2018; Liu et al., 2021b; Yu et al., 2021) methods are developing fast in recent years, which aim to measure heart activity remotely without any contact.

In earlier studies of facial rPPG measurement, most methods analyze subtle color changes on facial regions of interest (ROI) with classical signal processing approaches (Li et al., 2014; Magdalena Nowara et al., 2018; Poh et al., 2010a, 2010b; Tulyakov et al., 2016; Verkruysse et al., 2008). Besides, there are a few color subspace transformation methods (De Haan & Jeanne, 2013; Wang et al., 2016) which utilize all skin pixels for rPPG measurement. Based on the prior knowledge from traditional methods, a few learning based approaches (Hsu et al., 2017; Niu et al., 2018, 2019a; Qiu et al., 2018) are designed as non-end-to-end fashions. ROI based preprocessed signal representations [e.g., time-frequency map (Hsu et al., 2017) and spatio-temporal map (Niu et al., 2018, 2019a)] are generated first, and then learnable models could capture rPPG features from these maps. However, these methods need the strict preprocessing procedure and neglect the global contextual clues outside the pre-defined ROIs. Meanwhile, more and more end-to-end deep learning based rPPG methods (Chen & McDuff, 2018; Liu et al., 2020; Špetlík et al., 2018; Yu et al., 2019a, 2019b) are developed, which treat facial video frames as input and predict rPPG and other physiological signals directly. However, pure end-to-end methods are easily influenced by the complex scenarios (e.g., with head movement and various illumination conditions) and rPPG-unrelated features can not be ruled out in learning, resulting in huge performance drops (Yu et al., 2020) in realistic datasets [e.g., VIPL-HR (Niu et al., 2019a)].

Recently, due to its excellent long-range attentional modeling capacities in solving sequence-to-sequence issues,



**Fig. 1** The trajectories of rPPG signals around  $t_1$ ,  $t_2$ , and  $t_3$  share similar properties (e.g., trends with rising edge first then falling edge later, and relatively high magnitudes) induced by skin color changes. It inspires the long-range spatio-temporal attention (e.g., blue tube around  $t_1$  interacted with red tubes from intra- and inter-frames) according to their local temporal difference features for quasi-periodic rPPG enhancement. Here ‘tube’ indicates the same regions across short-time consecutive frames (Color figure online)

transformer (Han et al., 2022; Lin et al., 2022) has been successfully applied in many artificial intelligence tasks such as natural language processing (NLP) (Vaswani et al., 2017), image (Dosovitskiy et al., 2020) and video (Bertasius et al., 2021) analysis. Similarly, rPPG measurement from facial videos can be treated as a video sequence to signal sequence problem, where the long-range contextual clues should be exploited for semantic modeling. As shown in Fig. 1, rPPG clues from different skin regions and temporal locations (e.g., signal trajectories around  $t_1$ ,  $t_2$ , and  $t_3$ ) share similar properties (e.g., trends with rising edge first then falling edge later and relative high magnitudes), which can be utilized for long-range feature modeling and enhancement. However, different from the most video tasks aiming at semantic motion representation, facial rPPG measurement focuses on capturing subtle skin color changes, which makes it challenging for global spatio-temporal perception. Besides, the rPPG measurement task usually relies on periodic hidden visual dynamics, and the existing deep end-to-end models are weak in representing such clues. Furthermore, video-based rPPG measurement is usually a long-time monitoring task, and it is challenging to design and train transformers with long video sequence inputs.

Motivated by the discussions above, we propose two end-to-end video transformer architectures, namely PhysFormer and PhysFormer++, for remote physiological measurement. On the one hand, the cascaded temporal difference transformer blocks in PhysFormer benefit the rPPG feature enhancement via global spatio-temporal attention based on the fine-grained temporal skin color differences. Furthermore, the two-pathway SlowFast temporal difference transformer based PhysFormer++ with periodic- and cross-attention is able to efficiently capture the temporal contextual and periodic rPPG clues from facial videos. On the other hand, to alleviate the interference-induced overfitting issue and complement the weak temporal supervision signals, elaborate supervision in frequency domain is designed, which helps the PhysFormer family learn more intrinsic rPPG-aware features.

This paper is an extended version of our prior work (Yu et al., 2022) accepted by CVPR 2022. The main differences with the conference version are as follows: (1) besides the temporal difference transformer based PhysFormer, we propose the novel SlowFast video transformer architecture PhysFormer++ for rPPG measurement task; (2) based on the temporal difference transformer, the temporal difference periodic transformer and temporal difference cross-attention transformer are proposed to enhance the rPPG periodic perception and cross-tempo rPPG dynamics, respectively; (3) a detailed overview about the traditional, non-end-to-end learning based, and end-to-end learning based rPPG measurement methods is discussed in the related work; (4) more elaborate experimental results, visualization, and efficiency

analysis are given for the PhysFormer family. To sum up, the main contributions of this paper are listed:

- We propose the PhysFormer family, i.e., PhysFormer and PhysFormer++, which mainly consists of a powerful video temporal difference transformer backbone. To our best knowledge, it is the first time to explore the long-range spatio-temporal relationship for reliable rPPG measurement. Besides, the proposed temporal difference transformer is potential for broader fine-grained or periodic video understanding tasks in computer vision (e.g., video action recognition and repetition counting) due to its excellent spatio-temporal representation capacity with local temporal difference description and global spatio-temporal modeling.
- We propose the two-pathway SlowFast architecture for PhysFormer++ to efficiently leverage both fine-grained and semantic tempo rPPG clues. Specifically, the temporal difference periodic and cross-attention transformers are respectively designed for the Slow and Fast pathways to enhance the representation capacity of the periodic rPPG dynamics.
- We propose an elaborate recipe to supervise PhysFormer with label distribution learning and curriculum learning guided dynamic loss in frequency domain to learn efficiently and alleviate overfitting. Such curriculum learning guided dynamic strategy could benefit not only the rPPG measurement task but also general deep learning tasks such as multi-task learning and multi-loss adjusting.
- We conduct intra- and cross-dataset testings and show that the proposed PhysFormer achieves superior or on par state-of-the-art performance without pretraining on large-scale datasets like ImageNet-21K.

In the rest of the paper, Sect. 2 provides the related work about rPPG measurement and vision transformer. Section 3 first introduces the detailed architectures of the PhysFormer and PhysFormer++, and then formulates the label distribution learning and curriculum learning guided dynamic supervision for rPPG measurement. Section 4 introduces the four rPPG benchmark datasets and evaluation metrics, and provides rigorous ablation studies, visualizations and evaluates the performance of the proposed models. Finally, a conclusion is given in Sect. 5.

## 2 Related Work

In this section, we provide a brief discussion of the related facial rPPG measurement approaches. As shown in Table 1, these approaches can be generally categorized into traditional, non-end-to-end learning, and end-to-end learning

based methods. We also briefly review the transformer architectures for vision tasks.

### 2.1 rPPG Measurement

*Traditional Approaches* An early study of rPPG-based physiological measurement was reported in Verkruyse et al. (2008). Plenty of traditional hand-crafted approaches have been developed in this field since then. Compared with coarsely averaging arbitrary color channel from the detected full face region, selective merging information from different color channels (Poh et al., 2010a, 2010b) from different ROIs (Lam & Kuno, 2015; Li et al., 2014) with adaptive temporal filtering (Li et al., 2014) are proven to be more efficient for subtle rPPG signal recovery. To improve the signal-to-noise rate of the recovered rPPG signals, several signal decomposition methods such as independent component analysis (ICA) (Lam & Kuno, 2015; Poh et al., 2010a, 2010b) and matrix completion (Tulyakov et al., 2016) are also proposed. To alleviate the impacts of the skin tone and head motion, several color space projection [e.g., chrominance subspace (De Haan & Jeanne, 2013) and skin-orthogonal space (Wang et al., 2016)] methods are developed. Despite remarkable early-stage progresses, these approaches have the following limitations: (1) they require empirical knowledge to design the components (e.g., hyperparameter in signal processing filtering); (2) there is a lack of supervised learning models to counter data variations, especially in challenging environments with serious interference.

*Non-End-to-End Learning Approaches* In recent years, deep learning based approaches dominate the field of rPPG measurement due to the strong spatio-temporal representation capabilities. One representative framework is to learn robust rPPG features from the facial ROI-based spatio-temporal signal map (STmap). STmap (Niu et al., 2018, 2019b) or its variants [e.g., multisacle STmap (Lu et al., 2021; Niu et al., 2020) and chrominance STmap (Lu & Han, 2021)] are first extracted from predefined facial ROIs on different color spaces, and then classical convolutional neural network (CNN) [e.g., ResNet (He et al., 2016)] and recurrent neural network (RNN) [e.g., GRU (Cho et al., 2014)] are cascaded for rPPG feature representation. The STmap-based non-end-to-end learning framework focuses on learning an underlying mapping from the input feature maps to the target rPPG signals. With dense raw rPPG information and less irrelevant elements (e.g., face-shape attributes), these methods usually converge faster and achieve reasonable performance against head movement but need explicit and exhaustive preprocessings.

*End-to-End Learning Approaches* Besides learning upon handcrafted STmaps, end-to-end learning from facial sequence directly is also favorite. Both spatial 2DCNN networks

**Table 1** Summary of the representative rPPG measurement methods in terms of traditional, non-end-to-end learning, and end-to-end learning categories

| Method                               | Venue                    | Feature   | Backbone                      | Loss Function   |
|--------------------------------------|--------------------------|---|-------------------------------|---|
| Traditional                          |                          |   |                               |   |
| Poh2010 (Poh et al., 2010a)          | IEEE Trans. Biomed. Eng. | ROI selection, ICA decomposition  | –                             | –   |
| CHROM (De Haan & Jeanne, 2013)       | IEEE Trans. Biomed. Eng. | Mapping on the color difference (chrominance) subspace  | –                             | –   |
| Li2014 (Li et al., 2014)             | CVPR                     | ROI selection, tracking, illumination rectification, non-rigid motion elimination   | –                             | –   |
| RandomPatch (Lam & Kuno, 2015)       | ICCV                     | ICA decomposition on random patches, majority voting  | –                             | –   |
| Tulyakov2016 (Tulyakov et al., 2016) | CVPR                     | Chrominance features from multiple ROIs, low-rank factorization via self-adaptive matrix completion                         | –                             | –   |
| POS (Wang et al., 2016)              | IEEE Trans. Biomed. Eng. | Mapping on the projection plane orthogonal to the skin tone   | –                             | –   |
| SynRhythm (Niu et al., 2018)         | ICPR                     | spatio-temporal map representation, pretrained on Synthetic Rhythms   | ResNet18                      | L1 regression loss  |
| EVM-CNN (Qiu et al., 2018)           | IEEE TMM                 | ROI tracking, feature image via spatial decomposition and temporal filtering  | Shallow CNN                   | L2 regression loss  |
| RhythmNet (Niu et al., 2019a)        | IEEE TIP                 | spatio-temporal map in YUV color space, temporal contextual modeling via RNN  | ResNet18+GRU                  | L1 regression loss, temporal smooth loss                          |
| ST-Attention (Niu et al., 2019b)     | IEEE FG                  | spatio-temporal map in YUV color space, channel and spatio-temporal attention, temporal augmentation                        | ResNet18                      | L1 regression loss  |
| NAS-HR (Lu & Han, 2021)              | VRIH                     | Spatial-temporal map with POS signals, NAS for efficient architecture search  | lightweight NAS               | L1 regression loss  |
| CVD (Niu et al., 2020)               | ECCV                     | Multi-scale spatio-temporal map on both RGB and YUV color spaces, cross-verified feature disentangling, multi-task learning | ResNet18 with shallow decoder | L1 regression loss, CE loss, NegPearson loss, reconstruction loss |
| Dual-GAN (Lu et al., 2021)           | CVPR                     | spatio-temporal map, dual GAN for BVP signal and noise modeling, respectively   | Customized CNN                | L1 regression loss, CE loss, NegPearson loss, GAN loss            |

Table 1 continued

| Method                               | Venue    | Feature   | Backbone                              | Loss Function  |
|--------------------------------------|----------|---|---------------------------------------|--|
| End-to-end learning                  |          |   |                                       |  |
| DeepPhys (Chen & McDuff, 2018)       | ECCV     | Motion representation with normalized frame difference, attention mechanism using appearance to guide motion        | Two-branch CNN                        | MSE loss   |
| HR-CNN (Špetlík et al., 2018)        | BMVC     | Two-stage CNN to measure the rPPG signals first, and then estimate the HR value                                     | Shallow CNN                           | SNR loss, L1 regression loss                                 |
| PhysNet (Yu et al., 2019a)           | BMVC     | End-to-end spatio-temporal networks   | 3DCNN                                 | NegPearson loss  |
| rPPGNet (Yu et al., 2019b)           | ICCV     | Two-stage framework to enhance video quality first, and then extract rPPG signals                                   | 3DCNN                                 | NegPearson loss, reconstruction loss, skin segmentation loss |
| AutoHR (Yu et al., 2020)             | IEEE SPL | NAS with temporal difference convolution, spatio-temporal augmentation  | NAS 3DCNN                             | NegPearson loss, CE loss                                     |
| TS-CAN (Liu et al., 2020)            | NeuIPS   | Multi-task temporal shift convolutional attention networks, mobile-level real-time rPPG and respiratory measurement | Two-branch temporal shift CNN         | MSE loss on pulse and respiration                            |
| EfficientPhys (Liu et al., 2023)     | arXiv    | Temporal difference normalization, self-attention-shifted networks  | Temporal shift Swin Transformer       | MSE loss   |
| PhysFormer (Ours)PhysFormer++ (Ours) | 2022     | Temporal difference transformers supervised by the label distribution learning and curriculum learning strategy     | Temporal difference video transformer | NegPearson loss, CE loss, label distribution loss            |

(Chen & McDuff, 2018; Špetlík et al., 2018) and spatio-temporal models (Gideon & Stent, 2021; Liu et al., 2020, 2023; Nowara et al., 2021; Yu et al., 2019a, 2019b, 2020) are developed for rPPG feature representation. Yu et al. (2019a) investigates the recurrent methods (PhysNet-LSTM, PhysNet-ConvLSTM) for rPPG measurement. However, such CNN+LSTM based architectures are good at long-range sequential modeling via LSTM but fail to explore long-range intra-frame spatial relationship using CNN with local convolutions. In contrast, with spatial transformer backbone and temporal shift module, EfficientPhys (Liu et al., 2023) is able to explore long-range spatial but only short-term temporal relationship. In other words, existing end-to-end methods only consider the spatio-temporal rPPG features from local neighbors and adjacent frames but neglect the long-range relationship among quasi-periodic rPPG features.

Compared with the non-end-to-end learning based methods, end-to-end approaches are less dependent on task-related prior knowledge and handcrafted engineering (e.g., STmap generation) but rely on diverse and large-scale data to alleviate the problem of overfitting. To enhance the long-range contextual spatio-temporal representation capacities and alleviate the data-hungry requirement of the deep rPPG models, we propose the PhysFormer and PhysFormer++ architectures, which can be easily trained from scratch on rPPG datasets with the elaborate supervision recipe.

## 2.2 Transformer for Vision Tasks

Due to the powerful self-attention based long-range modeling capacity, transformer (Lin et al., 2022; Vaswani et al., 2017) has been successfully applied in the field of NLP to model the contextual relationship for sequential data. Then vision transformer (ViT) (Dosovitskiy et al., 2020) is proposed recently by feeding transformer with sequences of image patches for image classification. Many other ViT variants (Chen et al., 2021a; Ding et al., 2021; Han et al., 2022, 2021; Khan et al., 2021; Liu et al., 2021d; Touvron et al., 2021; Wang et al., 2021b; Yuan et al., 2021) are proposed from then, which achieve promising performance compared with its counterpart CNNs for image analysis tasks (Carion et al., 2020; He et al., 2021; Zheng et al., 2020). Recently, some works introduce vision transformer for video understanding tasks such as action recognition (Arnab et al., 2021; Bertasius et al., 2021; Bulat et al., 2021; Fan et al., 2021; Girdhar et al., 2018; Liu et al., 2021e; Neimark et al., 2021), action detection (Liu et al., 2021c; Wang et al., 2021a; Xu et al., 2021; Zhao et al., 2021), video super-resolution (Cao et al., 2021), video inpainting (Liu et al., 2021a; Zeng et al., 2020), and 3D animation (Chen et al., 2021b, 2021c). Some works (Girdhar et al., 2018; Neimark et al., 2021) conduct temporal contextual modeling with transformer based on single-frame features from pretrained 2D networks, while other works (Arnab et

al., 2021; Bertasius et al., 2021; Bulat et al., 2021; Fan et al., 2021; Liu et al., 2021e) mine the spatio-temporal attentions via video transformer directly. Most of these works are incompatible for long-video-sequence ( $> 150$  frames) signal regression task. There are two related works (Liu et al., 2023; Yu et al., 2022) using ViT for rPPG feature representation. TransRPPG (Yu et al., 2022) extracts rPPG features from the preprocessed signal maps via ViT for face 3D mask presentation attack detection (Yu et al., 2022). Based on the temporal shift networks (Lin et al., 2018; Liu et al., 2020), EfficientPhys-T (Liu et al., 2023) adds several Swin Transformer (Liu et al., 2021d) layers for global spatial attention. Different from these two works, the proposed PhysFormer and PhysFormer++ are end-to-end video transformers, which are able to capture long-range spatio-temporal attentional rPPG features from facial video directly.

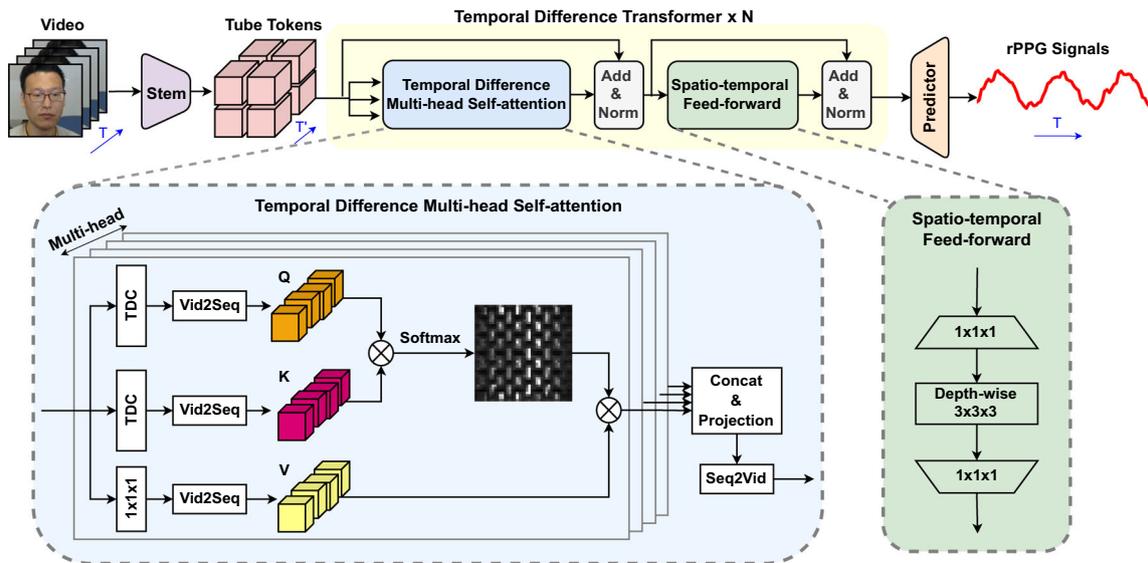
## 3 Methodology

We will first introduce the architecture of PhysFormer and PhysFormer++ in Sects. 3.1 and 3.2, respectively. Then we will introduce label distribution learning for rPPG measurement in Sect. 3.3, and at last present the curriculum learning guided dynamic supervision in Sect. 3.4.

### 3.1 PhysFormer

As illustrated in Fig. 2, PhysFormer consists of a shallow stem  $\mathbf{E}_{\text{stem}}$ , a tube tokenizer  $\mathbf{E}_{\text{tube}}$ ,  $N$  temporal difference transformer blocks  $\mathbf{E}_{\text{trans}}^i$  ( $i = 1, \dots, N$ ) and a rPPG predictor head. Inspired by the study in Xiao et al. (2021), we adopt a shallow stem to extract coarse local spatio-temporal features, which benefits the fast convergence and clearer subsequent global self-attention. Specifically, the stem is formed by three convolutional blocks with kernel size  $(1 \times 5 \times 5)$ ,  $(3 \times 3 \times 3)$  and  $(3 \times 3 \times 3)$ , respectively. Each convolution operator is cascaded with a batch normalization (BN), ReLU and MaxPool. The pooling layer only halves the spatial dimension. Therefore, given an RGB facial video input  $X \in \mathbb{R}^{3 \times T \times H \times W}$ , the stem output  $X_{\text{stem}} = \mathbf{E}_{\text{stem}}(X)$ , where  $X_{\text{stem}} \in \mathbb{R}^{D \times T \times H/8 \times W/8}$ , and  $D, T, W, H$  indicate channel, sequence length, width, height, respectively. Then  $X_{\text{stem}}$  would be partitioned into spatio-temporal tube tokens  $X_{\text{tube}} \in \mathbb{R}^{D \times T' \times H' \times W'}$  via the tube tokenizer  $\mathbf{E}_{\text{tube}}$ . Subsequently, the tube tokens will be forwarded with  $N$  temporal difference transformer blocks and obtain the global-local refined rPPG features  $X_{\text{trans}}$ , which has the same dimensions with  $X_{\text{tube}}$ . Finally, the rPPG predictor head temporally upsamples, spatially averages, and projects the features  $X_{\text{trans}}$  to 1D signal  $Y \in \mathbb{R}^T$ .

**Tube Tokenization** Here the coarse feature  $X_{\text{stem}}$  would be partitioned into non-overlapping tube tokens via  $\mathbf{E}_{\text{tube}}(X_{\text{stem}})$ ,



**Fig. 2** Framework of the PhysFormer. It consists of a shallow stem, a tube tokenizer, several temporal difference transformers, and a rPPG predictor head. The temporal difference transformer is formed from the Temporal Difference Multi-head Self-attention (TD-MHSA) and

Spatio-temporal Feed-forward (ST-FF) modules, which enhances the global and local spatio-temporal representation, respectively. ‘TDC’ is short for the temporal difference convolution (Yu et al., 2020, 2022)

which aggregates the spatio-temporal neighbor semantics within the tube region and reduces computational costs for the subsequent transformers. Specifically, the token tokenizer consists of a learnable 3D convolution with the same kernel size and stride (non-overlapping setting) as the targeted tube size  $T_s \times H_s \times W_s$ . Thus, the expected tube token map  $X_{tube} \in \mathbb{R}^{D \times T' \times H' \times W'}$  has length, height and width

$$T' = \left\lfloor \frac{T}{T_s} \right\rfloor, H' = \left\lfloor \frac{H/8}{H_s} \right\rfloor, W' = \left\lfloor \frac{W/8}{W_s} \right\rfloor. \tag{1}$$

Please note that there are no positional embeddings after the tube tokenization as the stem with cascaded convolutions and poolings at early stage already captures relative spatio-temporal positional information (Hassani et al., 2021).

*Temporal Difference Multi-head Self-Attention (TD-MHSA)*  
In self-attention mechanism (Dosovitskiy et al., 2020; Vaswani et al., 2017), the relationship between the tokens is modeled by the similarity between the projected query-key pairs, yielding the attention score. Instead of point-wise linear projection, we utilize temporal difference convolution (TDC) (Yu et al., 2020, 2022) for query ( $Q$ ) and key ( $K$ ) projection, which could capture fine-grained local temporal difference features for subtle color change description. TDC with learnable  $w$  can be formulated as

$$TDC(x) = \underbrace{\sum_{p_n \in \mathcal{R}} w(p_n) \cdot x(p_0 + p_n)}_{\text{vanilla 3D convolution}} + \theta \cdot \underbrace{(-x(p_0) \cdot \sum_{p_n \in \mathcal{R}'} w(p_n))}_{\text{temporal difference term}} \tag{2}$$

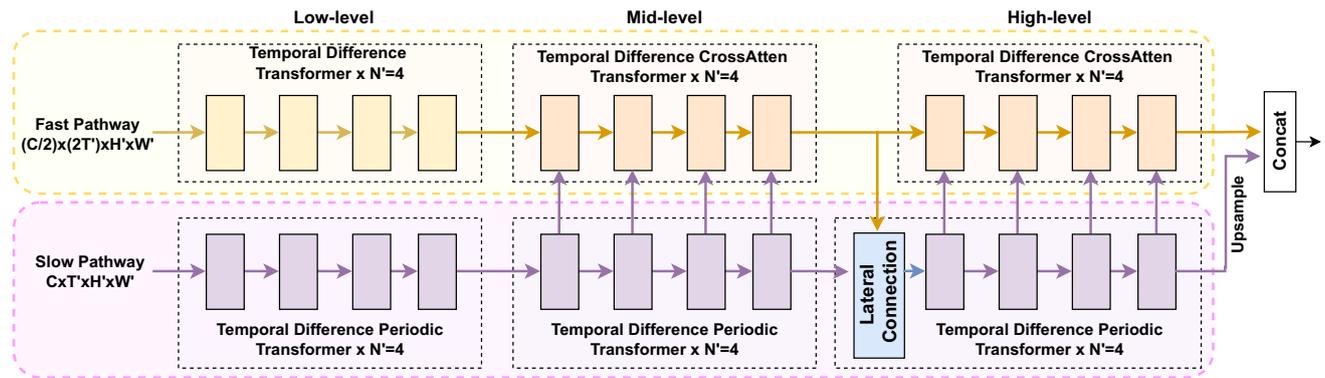
where  $p_0 = (0, 0, 0)$  indicates the current spatio-temporal location.  $\mathcal{R} = \{(-1, -1, -1), (-1, -1, 0), \dots, (0, 1, 1), (1, 1, 1)\}$  indicates the sampled local  $(3 \times 3 \times 3)$  spatio-temporal receptive field cube for 3D convolution in both current ( $t_0$ ) and adjacent time steps ( $t_{-1}$  and  $t_1$ ), while  $\mathcal{R}'$  only indicates the local spatial regions in the adjacent time steps ( $t-1$  and  $t1$ ). The hyperparameter  $\theta \in [0, 1]$  tradeoffs the contribution of temporal difference. The higher value of  $\theta$  means the more importance of temporal difference information (e.g., trends of the skin color changes). Specially, TDC degrades to vanilla 3D convolution when  $\theta = 0$ . Then query and key are projected via unshared TDC and BN as

$$Q = \text{BN}(\text{TDC}(X_{tube})), K = \text{BN}(\text{TDC}(X_{tube})). \tag{3}$$

For the value ( $V$ ) projection, point-wise linear projection without BN is utilized. Then  $Q, K, V \in \mathbb{R}^{D \times T' \times H' \times W'}$  are flattened into sequence, and separated into  $h$  heads ( $D_h = D/h$  for each head). For the  $i$ -th head ( $i \leq h$ ), the self-attention (SA) can be formulated

$$SA_i = \text{Softmax}(Q_i K_i^T / \tau) V_i, \tag{4}$$

where  $\tau$  controls the sparsity. We find that the default setting  $\tau = \sqrt{D_h}$  in Dosovitskiy et al. (2020); Vaswani et al. (2017) performs poorly for rPPG measurement. According to the periodicity of rPPG features, we use a smaller  $\tau$  value to obtain sparser attention activation. The corresponding study can be found in Table 8. The output of TD-MHSA is the concatenation of SA from all heads and then with a linear



**Fig. 3** Framework of the PhysFormer++ with two-stream SlowFast pathways. Different from the PhysFormer using only slow pathway, the PhysFormer++ extracts and fuses attentional features from slow and fast pathways. Moreover, temporal difference periodic transformer

blocks are used in the slow pathway. The information flow between two pathways interacts via temporal difference cross-attention transformer blocks and lateral connection

projection  $U \in \mathbb{R}^{D \times D}$

$$\text{TD-MHSA} = \text{Concat}(\text{SA}_1; \text{SA}_2; \dots; \text{SA}_h)U. \quad (5)$$

As illustrated in Fig. 2, residual connection and layer normalization (LN) would be conducted after TD-MHSA.

**Spatio-Temporal Feed-Forward (ST-FF)** The vanilla feed-forward network consists of two linear transformation layers, where the hidden dimension  $D'$  between two layers is expanded to learn a richer feature representation. In contrast, we introduce a depthwise 3D convolution (with BN and non-linear activation) between these two layers with extra slight computational cost but remarkable performance improvement. The benefits are twofold: (1) as a complementation of TD-MHSA, ST-FF could refine the local inconsistency and parts of noisy features; (2) richer locality provides TD-MHSA sufficient relative position cues.

### 3.2 PhysFormer++

In the PhysFormer, the temporal length  $T_s$  of the tube token map is fixed. However, the fixed value of  $T_s$  might be sub-optimal for robust rPPG feature representation as the larger  $T_s$  reduces the temporal redundancy but loses the fine-grained temporal clues, and vice versa for the smaller  $T_s$ . To alleviate this issue, we design the temporal enhanced version PhysFormer++ (see Fig. 3) consisting of two-stream SlowFast pathways with large and small  $T_s$ , respectively. Similar to the SlowFast concept in Feichtenhofer et al. (2018) and Kazakos et al. (2021), the Slow pathway has high channel capacity with low framerates, and reduces the temporal redundancy. In contrast, the Fast pathway operates at a fine-grained temporal resolution with high framerates. Furthermore, two novel transformer blocks, *temporal difference periodic transformer* and *temporal difference cross-attention transformer*,

are designed for the slow and fast pathway, respectively. The former one encodes contextual rPPG periodicity clues for the slow pathway while the latter one introduces efficient SlowFast interactive attentions for the fast pathway. The SlowFast architecture is able to adaptively mine richer temporally rPPG contexts for robust rPPG measurement.

As illustrated in Fig. 3 and detailed architecture in Table 2, different from the PhysFormer using a single tube tokenizer, two tube tokenizers  $\mathbf{E}_{\text{tube}}^{\text{fast}}$  and  $\mathbf{E}_{\text{tube}}^{\text{slow}}$  are adopted in PhysFormer++ to form the spatio-temporal tube tokens  $X_{\text{tube}}^{\text{fast}} \in \mathbb{R}^{D^{\text{fast}} \times T^{\text{fast}} \times H' \times W'}$  and  $X_{\text{tube}}^{\text{slow}} \in \mathbb{R}^{D^{\text{slow}} \times T^{\text{slow}} \times H' \times W'}$ , respectively. Default settings  $D^{\text{slow}} = D = 2D^{\text{fast}}$  and  $T^{\text{fast}} = 2T' = 2T^{\text{slow}}$  are used for computational tradeoff. Here we set temporal scale to two by considering that there are many low-framerate videos in the VIPL-HR dataset (Niu et al., 2019a). Too higher scales would result in the pulse rhythm incompleteness/artifacts for high HR values (e.g., >120 bpm). We will investigate more scales for higher framerate videos in the future. Subsequently, the tube tokens from the slow pathway will be forwarded with  $N = 3N'$  temporal difference periodic transformer blocks while tube tokens from the fast pathway will pass  $N'$  temporal difference transformer and  $2N'$  temporal difference cross-attention transformer blocks. Specifically, the feature interactions between SlowFast pathways are in two folds: (1) all semantic mid- and high-level features from the slow path are cross-attentive with those from the fast path; and (2) the last mid-level features from two pathways  $X_{\text{tube}}^{\text{fast-mid}}$ ,  $X_{\text{tube}}^{\text{slow-mid}}$  are lateral connected and then aggregated for the high-level propagation in the slow pathway. The lateral connection and aggregation can be formulated as

$$X_{\text{tube}}^{\text{slow-mid}} = \text{Conv2} \left( \text{Concat} \left( X_{\text{tube}}^{\text{slow-mid}}, \text{Conv1} \left( X_{\text{tube}}^{\text{fast-mid}} \right) \right) \right), \quad (6)$$

**Table 2** Architectures of PhysFormer++

| Stage             | Slow pathway   | Fast pathway  | Output $C \times T \times H \times W$  |
|-------------------|--|---|--|
| Input             | –  | –   | $3 \times 160 \times 128 \times 128$   |
| Stem              | $\left[ \begin{array}{l} 1 \times 5 \times 5 \text{ Conv, } 24 \\ 1 \times 2 \times 2 \text{ MaxPool} \\ 3 \times 3 \times 3 \text{ Conv, } 48 \\ 1 \times 2 \times 2 \text{ MaxPool} \\ 3 \times 3 \times 3 \text{ Conv, } 96 \\ 1 \times 2 \times 2 \text{ MaxPool} \end{array} \right]$ |   | $96 \times 160 \times 16 \times 16$  |
| Tokenizer         | $4 \times 4 \times 4 \text{ Conv, } 96$<br>stride = $4 \times 4 \times 4$  | $2 \times 4 \times 4 \text{ Conv, } 48$<br>stride = $2 \times 4 \times 4$ | Slow: $96 \times 40 \times 4 \times 4$<br>Fast: $48 \times 80 \times 4 \times 4$ |
| Low-level blocks  | Temporal difference periodic transformer $\times 4$  | Temporal difference transformer $\times 4$                                | Slow: $96 \times 40 \times 4 \times 4$<br>Fast: $48 \times 80 \times 4 \times 4$ |
| Mid-level blocks  | Temporal difference periodic transformer $\times 4$  | Temporal difference CrossAtten transformer $\times 4$                     | Slow: $96 \times 40 \times 4 \times 4$<br>Fast: $48 \times 80 \times 4 \times 4$ |
| High-level blocks | Lateral connection $\times 1$ temporal difference periodic transformer $\times 4$  | Temporal difference CrossAtten transformer $\times 4$                     | Slow: $96 \times 40 \times 4 \times 4$<br>Fast: $48 \times 80 \times 4 \times 4$ |
| Pre-fusion        | $\left[ \begin{array}{l} 2 \times 1 \times 1 \text{ Upsample} \\ 3 \times 1 \times 1 \text{ Conv, } 96 \end{array} \right]$  | $3 \times 1 \times 1 \text{ Conv, } 48$                                   | Slow: $96 \times 80 \times 4 \times 4$<br>Fast: $48 \times 80 \times 4 \times 4$ |
| Fusion            | Concat (Slow, Fast), 144   |   | $144 \times 80 \times 4 \times 4$  |
| Predictor         | $\left[ \begin{array}{l} 2 \times 1 \times 1 \text{ Upsample} \\ 3 \times 1 \times 1 \text{ Conv, } 72 \\ 1 \times 4 \times 4 \text{ AvgPool} \\ 3 \times 1 \times 1 \text{ Conv, } 1 \end{array} \right]$   |   | $1 \times 160 \times 1 \times 1$   |

Inside the brackets are the filter sizes and feature dimensionalities. ‘Conv’ suggests the vanilla 3D convolution. All convolutional layers (except Tokenizer) are with stride=1 and are followed by a BN+ReLU layer while ‘MaxPool’ layers are with stride=1x2x2

where Conv1 is the temporal convolution with size =  $3 \times 1 \times 1$ , stride =  $2 \times 1 \times 1$ , padding =  $1 \times 0 \times 0$  while Conv2 denotes the point-wise convolution with  $D$  output channel. The lateral connection adaptively transfers the mid-level fine-grained rPPG clues from the Fast pathway to the Slow pathway, and provides complementary temporal details for the Slow pathway to alleviate information loss especially for high-HR scenarios (e.g., after exercise). Finally, the refined high-level rPPG features from fast and slow (upsampled) pathways are concatenated and forwarded the rPPG predictor head with temporally aggregation, upsampling, spatially averaging, and 1D signal  $\hat{Y} \in \mathbb{R}^T$  projection.

**Temporal Difference Multi-head Cross- and Self-Attention**  
 Compared with the slow pathway, the fast pathway has more fine-grained features but conducts inefficient and inaccurate self-attention due to the temporal redundancy/artifacts. To alleviate the weak self-attention issue in the fast pathway, we propose the temporal difference multi-head cross- and self-attention (TD-MHCSA) module, which could be cascaded with ST-FF module to form the temporal difference cross-attention transformer. With TD-MHCSA, the features in the fast pathway can not only be refined by its own self-attention but also the cross-attention between the SlowFast pathways.

The structure of the TD-MHCSA is illustrated in Fig. 4. The features from the fast pathway  $X_{tube}^{fast}$  are first projected to query and key via

$$Q^{fast} = \text{BN}(\text{TDC}(X_{tube}^{fast})), K^{fast} = \text{BN}(\text{TDC}(X_{tube}^{fast})). \quad (7)$$

For the value ( $V^{fast}$ ) projection, point-wise linear projection without BN is utilized. Then  $Q^{fast}, K^{fast}, V^{fast} \in \mathbb{R}^{D^{fast} \times T^{fast} \times H' \times W'}$  are flattened into sequence, and separated into  $h$  heads ( $D_h^{fast} = D^{fast}/h$  for each head). For the  $i$ -th head ( $i \leq h$ ), the self-attention can be formulated

$$SA_i^{fast} = \text{Softmax}(Q_i^{fast} K_i^{fastT} / \tau) V_i^{fast}. \quad (8)$$

Similarly, the features from the slow pathway  $X_{tube}^{slow}$  are projected to key  $K^{slow}$  via  $\text{BN}(\text{TDC}(X_{tube}^{slow}))$  as well as the value ( $V^{slow}$ ) projection using point-wise linear projection. Then  $K^{slow}, V^{slow} \in \mathbb{R}^{D^{slow} \times T^{slow} \times H' \times W'}$  are flattened into sequence, and separated into  $h$  heads. For the  $i$ -th head ( $i \leq h$ ), the cross-attention (CA) can be formulated as

$$CA_i = \text{Softmax}(Q_i^{fast} K_i^{slowT} / \tau) V_i^{slow}. \quad (9)$$

Thus, the combined cross- and self-attention (CSA) is formulated as  $CSA_i = CA_i + SA_i^{fast}$ . The output of TD-MHCSA is the concatenation of CSA from all heads and then with a linear projection  $U^{fast} \in \mathbb{R}^{D^{fast} \times D^{fast}}$ , which is formulated

$$\text{TD-MHCSA} = \text{Concat}(SCA_1; SCA_2; \dots; SCA_h) U^{fast}. \quad (10)$$

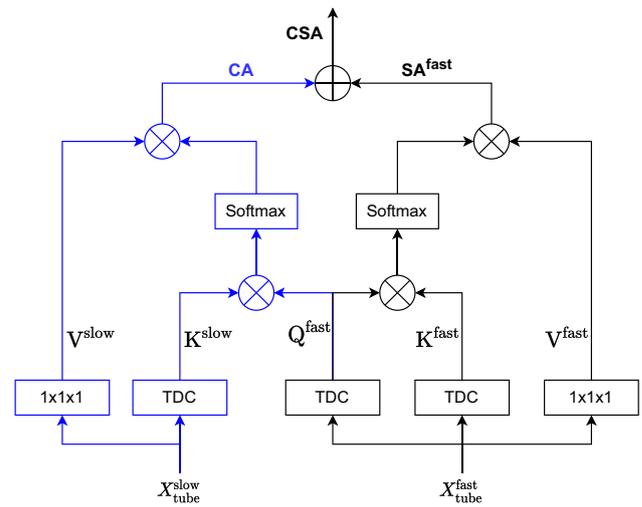


Fig. 4 Illustration of the temporal difference multi-head cross- and self-attention (TD-MHCSA) module

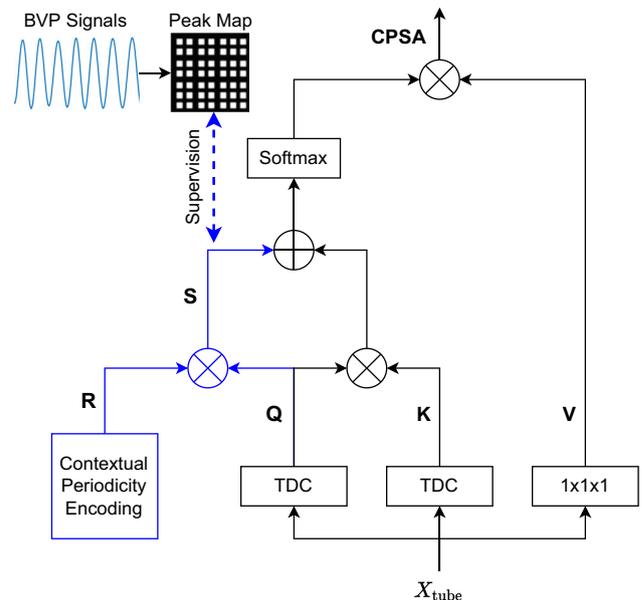


Fig. 5 Illustration of the temporal difference multi-head periodic- and self-attention (TD-MHPSA) module

Finally, residual connection and LN layer would be conducted after TD-MHCSA.

**Temporal Difference Multi-head Periodic- and Self-Attention**  
 Inspired by the music transformer (Huang et al., 2019) using relative attention (Shaw et al., 2018; Wu et al., 2021) to mine richer positional relationship (e.g., periodicity in music signals), we propose the temporal difference multi-head periodic- and self-attention (TD-MHPSA), which extends the TD-MHSA (in Sect. 3.1) with learnable rPPG-aware positional contextual periodicity representation. Specifically, as shown in Fig. 5, the learnable contextual periodicity encoding  $R \in \mathbb{R}^{T' \times H' \times W' \times T' \times H' \times W' \times D}$  contains the spatio-temporal

positional clues, and modulates the query vector  $Q$  into the periodic attention  $S = QR^T$ . In consideration of the multi-head  $h$  setting, for the  $i$ -th head, the joint contextual periodicity (CP) and self-attention (SA) can be formulated as

$$CPSA_i = \text{Softmax}((Q_i K_i^T + \lambda \cdot S_i)/\tau) V_i, \tag{11}$$

where  $\lambda$  tradeoffs the CP and SA. Here we follow the memory efficient implementation in Huang et al. (2019) for  $S$  calculation.

Despite richer positional periodicity clues, the predicted periodic attention  $S$  might be easily influenced by some rPPG-unrelated clues (e.g., light changes and dynamic noise). To alleviate this issue, we propose a periodicity constraint to supervise the periodic  $S$  representation. As shown in the top left of Fig. 5, the approximate peak map PM can be obtained via (1) first extracting the binary peak signal  $P \in \mathbb{R}^T$  from the ground truth BVP signal  $Y \in \mathbb{R}^T$  via

$$P_{t \in T} = \begin{cases} 1, & \text{if } Y_t \in \mathcal{R}_{peak}, \\ 0, & \text{if } Y_t \notin \mathcal{R}_{peak}, \end{cases} \tag{12}$$

where  $\mathcal{R}_{peak}$  denotes the 1D-region of peak locations; and then (2) calculating the auto-correlation of the peak signal  $P$  via  $PM = PP^T$ . Finally, the periodic-attention loss  $\mathcal{L}_{atten}$  can be calculated with the binary cross-entropy (BCE) loss between the adaptive-spatial-pooled periodic attention maps  $S' \in \mathbb{R}^{T' \times T'}$  (from each head and each TD-MHPSA module) and the subsampled binary peak maps  $PM' \in \mathbb{R}^{T' \times T'}$ . It can be formulated as

$$\mathcal{L}_{atten} = \frac{1}{h \times N} \sum_{i \in h, j \in N} \text{BCE}(S', PM'). \tag{13}$$

We also try supervision with  $L1$  regression loss instead of BCE loss but with poorer performance.

*Relationship Between PhysFormer and PhysFormer++* PhysFormer++ can be treated as an upgraded version of PhysFormer towards excellent performance while with more computational cost. With similar temporal difference transformers, PhysFormer can be seen as a slow-pathway version of PhysFormer++, which is more lightweight and efficient. In contrast, PhysFormer++ is designed based on a dual-pathway SlowFast architecture with complex cross-tempo interactions, which is more robust to head motions and less sensitive to the video framerate, but with heavier computational cost (see Table 12 for efficiency analysis).

### 3.3 Label Distribution Learning

Similar to the facial age estimation task (Gao et al., 2018; Geng et al., 2010) that faces at close ages look quite simi-

lar, facial rPPG signals with close HR values usually have similar periodicity. Inspired by this observation, instead of considering each facial video as an instance with one label (HR), we regard each facial video as an instance associated with a label distribution. The label distribution covers a certain number of class labels, representing the degree that each label describes the instance. Through this way, one facial video can contribute to both targeted HR value and its adjacent HRs.

To consider the similarity information among HR classes during the training stage, we model the rPPG-based HR estimation problem as a specific  $L$ -class multi-label classification problem, where  $L = 139$  in our case (each integer HR value within [42, 180] bpm as a class). A label distribution  $\mathbf{p} = \{p_1, p_2, \dots, p_L\} \in \mathbb{R}^L$  is assigned to each facial video  $X$ . It is assumed that each entry of  $\mathbf{p}$  is a real value in the range [0,1] such that  $\sum_{k=1}^L p_k = 1$ . We consider the Gaussian distribution function, centered at the ground truth HR label  $Y_{HR}$  with the standard deviation  $\sigma$ , to construct the corresponding label distribution  $\mathbf{p}$ .

$$p_k = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(k - (Y_{HR} - 41))^2}{2\sigma^2}\right). \tag{14}$$

The label distribution loss can be formulated as  $\mathcal{L}_{LD} = \text{KL}(\mathbf{p}, \text{Softmax}(\hat{\mathbf{p}}))$ , where divergence measure  $\text{KL}(\cdot)$  denotes the Kullback-Leibler (KL) divergence (Gao et al., 2016), and  $(\hat{\mathbf{p}})$  is the power spectral density (PSD) of predicted rPPG signals.

Please note that the previous work (Niu et al., 2017) also considers the distribution learning for HR estimation. However, it is totally different with our work: (1) the motivation in Niu et al. (2017) is to smooth the temporal HR outliers caused by facial movements across continuous video clips, while our work is more generic, aiming at efficient feature learning across adjacent labels under limited-scale training data; (2) the technique used in Niu et al. (2017) is after a post-HR-estimation for the handcrafted rPPG signals, while our work is to design a reasonable supervision signal  $\mathcal{L}_{LD}$  for the PhysFormer family.

### 3.4 Curriculum Learning Guided Dynamic Loss

Curriculum learning (Bengio et al., 2009), as a major machine learning regime with philosophy of easy-to-hard curriculum, is utilized to train PhysFormer. In the rPPG measurement task, the supervision signals from temporal domain [e.g., mean square error loss (Chen & McDuff, 2018), negative Pearson loss (Yu et al., 2019a,b)] and frequency domain [e.g., cross-entropy loss (Niu et al., 2020; Yu et al., 2020), signal-to-noise ratio loss (Špetlík et al., 2018)] provide different extents of constraints for model learning. The former one gives signal-trend-level constraints, which is straight-

forward for model convergence but overfitting after that. In contrast, the latter one with strong constraints on frequency domain enforces the model learning periodic features within target frequency bands, which is hard to converge well due to the realistic rPPG-irrelevant noise. Inspired by the curriculum learning, we propose the dynamic supervision to gradually enlarge the frequency constraints, which alleviates the overfitting issue and benefits the intrinsic rPPG-aware feature learning gradually. Specifically, exponential increment strategy is adopted, and comparison with other dynamic strategies (e.g., linear increment) will be shown in Table 11. The dynamic loss  $\mathcal{L}_{\text{overall}}$  can be formulated as

$$\mathcal{L}_{\text{overall}} = \underbrace{\alpha \cdot \mathcal{L}_{\text{time}}}_{\text{temporal}} + \underbrace{\beta \cdot (\mathcal{L}_{\text{CE}} + \mathcal{L}_{\text{LD}})}_{\text{frequency}} + \mathcal{L}_{\text{atten}}, \quad (15)$$

$$\beta = \beta_0 \cdot (\eta)^{(\text{Epoch}_{\text{current}} - 1) / \text{Epoch}_{\text{total}}},$$

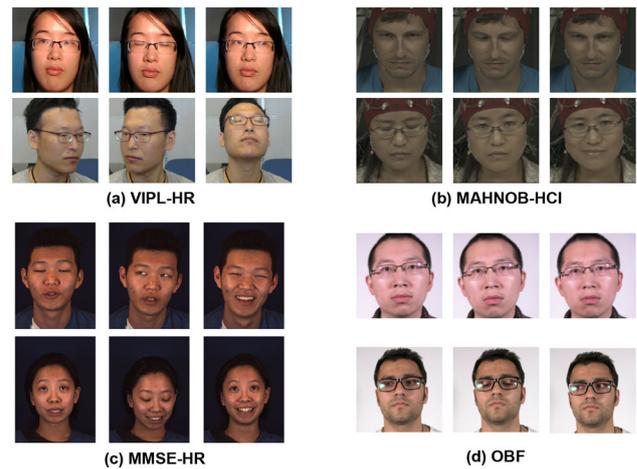
where hyperparameters  $\alpha$ ,  $\beta_0$  and  $\eta$  equal to 0.1, 1.0 and 5.0, respectively. Negative Pearson loss (2019a, 2019b) and frequency cross-entropy loss (Niu et al., 2020; Yu et al., 2020) are adopted as  $\mathcal{L}_{\text{time}}$  and  $\mathcal{L}_{\text{CE}}$ , respectively. With the dynamic supervision, PhysFormer and PhysFormer++ could perceive better signal trend at the beginning while such perfect warming up facilitates the gradually stronger frequency knowledge learning later.

## 4 Experimental Evaluation

In this section, experiments of rPPG-based physiological measurement for three types of physiological signals, i.e., heart rate (HR), heart rate variability (HRV), and respiration frequency (RF), are conducted on four benchmark datasets (VIPL-HR (Niu et al., 2019a), MAHNOB-HCI (Soleymani et al., 2012), MMSE-HR (Tulyakov et al., 2016), and OBF (Li et al., 2018)). Besides, comprehensive ablations about PhysFormer and PhysFormer++ are also investigated in the VIPL-HR dataset.

### 4.1 Datasets and Performance Metrics

**VIPL-HR** (Niu et al., 2019a) is a large-scale dataset for remote physiological measurement under less-constrained scenarios. It contains 2,378 RGB videos of 107 subjects recorded with different head movements, lighting conditions and acquisition devices. **MAHNOB-HCI** (Soleymani et al., 2012) is one of the most widely used benchmark for remote HR measurement evaluations. It includes 527 facial videos of with 61 fps framerate and  $780 \times 580$  resolution from 27 subjects. **MMSE-HR** (Tulyakov et al., 2016) is a dataset including 102 RGB videos from 40 subjects, and the raw resolution of each video is at  $1040 \times 1392$ . **OBF** (Li et al., 2018)



**Fig. 6** Example video frames from datasets **a** VIPL-HR (Niu et al., 2019a); **b** MAHNOB-HCI (Soleymani et al., 2012); **c** MMSE-HR (Tulyakov et al., 2016); and **d** OBF (Li et al., 2018)

is a high-quality dataset for remote physiological signal measurement. It contains 200 five-minute-long RGB videos with 60 fps framerate recorded from 100 healthy adults. The example video frames from these four rPPG datasets are illustrated in Fig. 6.

For MAHNOB-HCI, as there are no available BVP ground truth, we first smooth the sharp ECG signals (with 10-point averaging strategy) into pseudo BVP signals as ground truth. Specifically, to alleviate the incorrect synchronization between videos and ground truth signals in MAHNOB-HCI, OBF, and VIPL-HR datasets, we first extract the coarse green channel signals via averaging the segmented facial skin in each frame. Then, we calculate the cross-correlation between the coarse green rPPG signals and (pseudo) BVP signals, and use the maximum-correlation phase to calibrate/compensate the phase bias. Furthermore, we remove the samples with HR > 180 in the VIPL-HR and MMSE-HR datasets because the ground truths in these samples are unreliable due to poor contact of sensors (resulting in very noisy and fluctuated HRs).

In terms of evaluation metrics, average HR estimation task is evaluated on all four datasets while HRV and RF estimation tasks on high-quality OBF (Li et al., 2018) dataset. Specifically, we follow existing methods (Lu et al., 2021; Niu et al., 2020; Yu et al., 2019b) and report low frequency (LF), high frequency (HF), and LF/HF ratio for HRV and RF estimation. We report the most commonly used performance metrics for evaluation, including the standard deviation (SD), mean absolute error (MAE), root mean square error (RMSE), and Pearson's correlation coefficient ( $r$ ).

### 4.2 Implementation Details

Both PhysFormer and PhysFormer++ are implemented with Pytorch. For each video clip, the MTCNN face detector

**Table 3** Intra-dataset testing results on the VIPL-HR dataset

| Method                                   | SD ↓ (bpm)  | MAE ↓ (bpm) | RMSE ↓ (bpm) | $r$ ↑       |
|--|-------------|-------------|--------------|-------------|
| Tulyakov2016 (Tulyakov et al., 2016)▲    | 18.0        | 15.9        | 21.0         | 0.11        |
| POS (Wang et al., 2016)▲                 | 15.3        | 11.5        | 17.2         | 0.30        |
| CHROM (De Haan & Jeanne, 2013)▲          | 15.1        | 11.4        | 16.9         | 0.28        |
| RhythmNet (Niu et al., 2019a)◆           | 8.11        | 5.30        | 8.14         | 0.76        |
| ST-Attention (Niu et al., 2019b)◆        | 7.99        | 5.40        | 7.99         | 0.66        |
| NAS-HR (Lu & Han, 2021)◆                 | 8.10        | 5.12        | 8.01         | 0.79        |
| CVD (Niu et al., 2020)◆                  | 7.92        | 5.02        | 7.97         | 0.79        |
| Dual-GAN (Lu et al., 2021)◆              | <b>7.63</b> | <u>4.93</u> | <u>7.68</u>  | <b>0.81</b> |
| I3D (Carreira & Zisserman, 2017)★        | 15.9        | 12.0        | 15.9         | 0.07        |
| PhysNet (Yu et al., 2019a)★              | 14.9        | 10.8        | 14.8         | 0.20        |
| DeepPhys (Chen & McDuff, 2018)★          | 13.6        | 11.0        | 13.8         | 0.11        |
| VideoTransformer (Revanur et al., 2022)★ | 13.5        | 10.4        | 13.2         | 0.16        |
| AutoHR (Yu et al., 2020)★                | 8.48        | 5.68        | 8.68         | 0.72        |
| PhysFormer (Ours)★                       | 7.74        | 4.97        | 7.79         | 0.78        |
| PhysFormer++ (Ours)★                     | <u>7.65</u> | <b>4.88</b> | <b>7.62</b>  | <u>0.80</u> |

The symbols ▲, ◆ and ★ denote traditional, non-end-to-end learning based and end-to-end learning based methods, respectively. Best results are marked in bold and second best in underline

**Table 4** Intra-dataset results on the MAHNOB-HCI dataset

| Method                                | SD ↓ (bpm)  | MAE ↓ (bpm) | RMSE ↓ (bpm) | $r$ ↑       |
|---------------------------------------|-------------|-------------|--------------|-------------|
| Poh2010 (Poh et al., 2010a)▲          | 13.5        | –           | 13.6         | 0.36        |
| CHROM (De Haan & Jeanne, 2013)▲       | –           | 13.49       | 22.36        | 0.21        |
| Li2014 (Li et al., 2014)▲             | 6.88        | –           | 7.62         | 0.81        |
| Tulyakov2016 (Tulyakov et al., 2016)▲ | 5.81        | 4.96        | 6.23         | 0.83        |
| SynRhythm (Niu et al., 2018)◆         | 10.88       | –           | 11.08        | –           |
| RhythmNet (Niu et al., 2019a)◆        | 3.99        | –           | 3.99         | <b>0.87</b> |
| HR-CNN (Špetlík et al., 2018)★        | –           | 7.25        | 9.24         | 0.51        |
| rPPGNet (Yu et al., 2019b)★           | 7.82        | 5.51        | 7.82         | 0.78        |
| DeepPhys (Chen & McDuff, 2018)★       | –           | 4.57        | –            | –           |
| AutoHR (Yu et al., 2020)★             | 4.73        | 3.78        | 5.10         | <u>0.86</u> |
| Meta-rPPG (Lee et al., 2020)★         | 4.9         | <b>3.01</b> | <b>3.68</b>  | 0.85        |
| PhysFormer (Ours)★                    | <b>3.87</b> | 3.25        | 3.97         | <b>0.87</b> |
| PhysFormer++ (Ours)★                  | <u>3.90</u> | <u>3.23</u> | <u>3.88</u>  | <b>0.87</b> |

The symbols ▲, ◆ and ★ denote traditional, non-end-to-end learning based and end-to-end learning based methods, respectively. Best results are marked in bold and second best in underline

(Zhang et al., 2016) is used to crop the enlarged face area at the first frame and fix the region through the following frames. The videos in MAHNOB-HCI and OBF are down-sampled to 30 fps for efficiency. The numbers of temporal difference transformer blocks  $N = 12$ , transformer heads  $h = 4$ , channel dimension  $D = 96$ , hidden dimension in ST-FF  $D' = 144$  are used for PhysFormer while temporal difference coefficient  $\theta = 0.7$  and attention sparsity  $\tau = 2.0$  for TD-MHSA.  $\lambda = 0.5$  is utilized in the TD-MHPSA. The targeted tube size  $T_s \times H_s \times W_s$  equals to  $4 \times 4 \times 4$ . For the  $\mathcal{R}_{peak}$  calculation in Eq. (12), the function ‘findpeaks()’ in Matlab is used for BVP peak detection, and then the detected peak locations are extended with successive  $\pm 3$  neighbors.

In the training stage, we randomly sample RGB face clips with size  $160 \times 128 \times 128$  ( $T \times H \times W$ ) as model inputs. Random horizontal flipping and temporally up/down-sampling (Yu et al., 2020) are used for data augmentation. The PhysFormer is trained with Adam optimizer and the initial learning rate and weight decay are  $1e-4$  and  $5e-5$ , respectively. We cannot find obvious performance improvement using AdamW optimizer. We train models with 25 epochs with fixed setting  $\alpha = 0.1$  for temporal loss while exponentially increased parameter  $\beta \in [1, 5]$  for frequency losses. We set standard deviation  $\sigma = 1.0$  for label distribution learning. The batch size is 4 on one V100 GPU. In the testing stage, similar to Niu et al. (2019a), we uniformly

separate 30-second videos into three short clips with 10 seconds, and then video-level HR is calculated via averaging HRs from three short clips.

### 4.3 Intra-dataset Testing

In this subsection, two datasets (VIPL-HR and MAHNOB-HCI) are used for intra-dataset testing on HR estimation while the OBF dataset is used for intra-dataset HR, HRV and RF estimation.

**HR Estimation on VIPL-HR** Here we follow (Niu et al., 2019a) and use a subject-exclusive 5-fold cross-validation protocol on VIPL-HR. As shown in Table 3, all three traditional methods [Tulyakov2016 (Tulyakov et al., 2016), POS (Wang et al., 2016) and CHROM (De Haan & Jeanne, 2013)] perform poorly due to the complex scenarios (e.g., large head movement and various illumination) in the VIPL-HR dataset. In terms of deep learning based methods, the existing end-to-end learning based methods [e.g., PhysNet (Yu et al., 2019a), DeepPhys (Chen & McDuff, 2018), and AutoHR (Yu et al., 2020)] predict less reliable HR values with larger RMSE compared with non-end-to-end learning approaches [e.g., RhythmNet (Niu et al., 2019a), ST-Attention (Niu et al., 2019b), NAS-HR (Lu & Han, 2021), CVD Niu et al. (2020), and Dual-GAN (Lu et al., 2021)]. Such the large performance margin might be caused by the coarse and overfitted rPPG features extracted from the end-to-end models. In contrast, all five non-end-to-end methods first extract fine-grained signal maps from multiple facial ROIs, and then more dedicated rPPG clues would be extracted via the cascaded models. Without strict and heavy preprocessing procedure in Niu et al. (2019a, 2019b, 2020), Lu and Han (2021) and Lu et al. (2021), the proposed PhysFormer and PhysFormer++ can be trained from scratch on facial videos directly, and achieve better or on par performance with state-of-the-art non-end-to-end learning based method Dual-GAN (Lu et al., 2021). It indicates that PhysFormer and PhysFormer++ are able to learn the intrinsic and periodic rPPG-aware features automatically. It can be seen from Table 3 that the proposed PhysFormer family outperforms the VideoTransformer (Revanur et al., 2022) by a large margin, indicating the importance of local and global spatio-temporal physiological propagation.

In order to further check the correlations between the predicted HRs and the ground-truth HRs, we plot the HR estimation results against the ground truths in Fig. 7a. From the figure we can see that the predicted HRs from PhysFormer++ and the ground-truth HRs are well correlated in a wide range of HR from 47 to 147 bpm.

**HR Estimation on MAHNOB-HCI** For the HR estimation tasks on MAHNOB-HCI, similar to Yu et al. (2019b), subject-independent 9-fold cross-validation proto-

col is adopted. In consideration of the convergence difficulty due to the low illumination and high compression videos in MAHNOB-HCI, we finetune the VIPL-HR pretrained models on MAHNOB-HCI for further 15 epochs. The HR estimation results are shown in Table 4. The proposed PhysFormer and PhysFormer++ achieves the lowest SD (3.87 bpm) and highest  $r$  (0.87) among the traditional, non-end-to-end learning, and end-to-end learning methods, which indicates the reliability of the learned rPPG features from PhysFormer family under sufficient supervision. Our performance is on par with the latest end-to-end learning method Meta-rPPG (Lee et al., 2020) without transductive adaptation from target frames.

**HR, HRV and RF Estimation on OBF** Besides HR estimation, we also conduct experiments for three types of physiological signals, i.e., HR, RF, and HRV measurement on the OBF (Li et al., 2018) dataset. Following Niu et al. (2020) and Yu et al. (2019b), we use a 10-fold subject-exclusive protocol for all experiments. All the results are shown in Table 5. It is clear that the proposed PhysFormer and PhysFormer++ outperform the existing state-of-the-art traditional [ROI\_green (Li et al., 2018)), CHROM (De Haan & Jeanne, 2013), POS (Wang et al., 2016)] and end-to-end learning [rPPGNet (Yu et al., 2019b)] methods by a large margin on all evaluation metrics for HR, RF and all HRV features. The proposed PhysFormer and PhysFormer++ give more accurate estimation in terms of HR, RF, and LF/HF compared with the preprocessed signal map based non-end-to-end learning method CVD (Niu et al., 2020). These results indicate that PhysFormer family could not only handle the average HR estimation task but also give a promising prediction of the rPPG signal for RF measurement and HRV analysis, which shows its potential in many healthcare applications.

We also check the short-time HR estimation performance of the after exercising scenario on the OBF, in which the subject's HR decreases rapidly. Two examples are given in Fig. 7b. It can be seen that PhysFormer++ could follow the trend of HR changes well, which indicates the proposed model is robust in the significant HR changing scenarios. We further check the predicted rPPG signals of the PhysFormer++ from these two examples in Fig. 7c. From the results, we can see that the proposed method could give an accurate prediction of the interbeat intervals (IBIs), thus can give a robust estimation of RF and HRV features (Table 6).

### 4.4 Cross-dataset Testing

Besides of the intra-dataset testings on the VIPL-HR, MAHNOB-HCI, and OBF datasets, we also conduct cross-dataset testings on MMSE-HR (Tulyakov et al., 2016) following the protocol of Niu et al. (2019a). The models trained on VIPL-HR are directly tested on MMSE-HR. All

**Table 5** Performance comparison of HR and RF measurement as well as HRV analysis on the OBF dataset

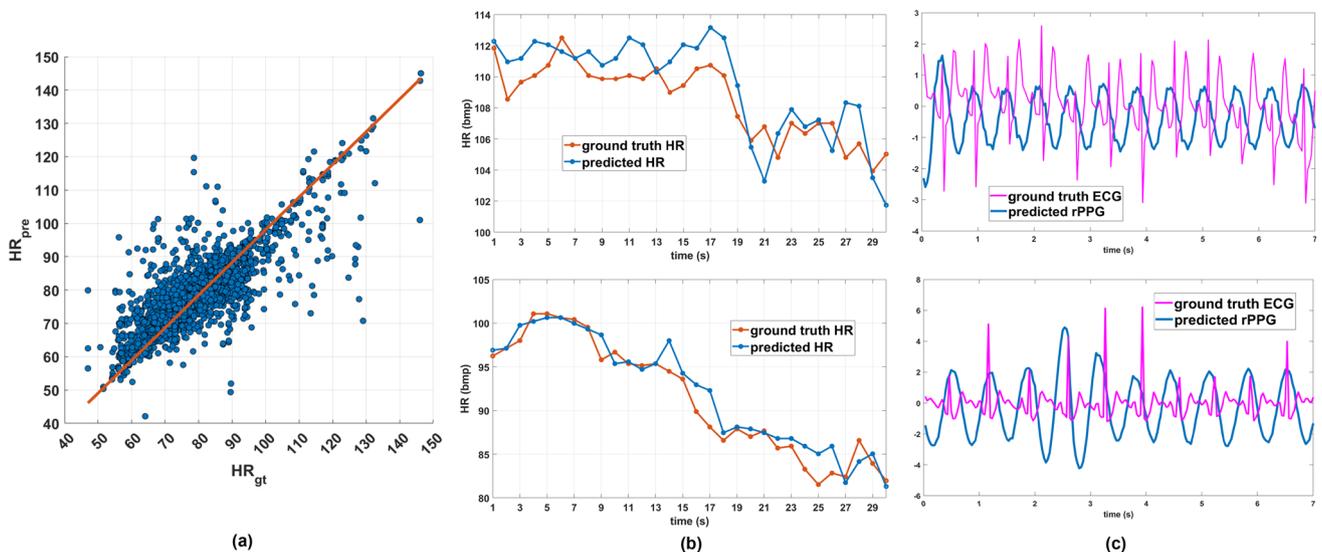
| Method                          | HR (bpm)     |              | RF (Hz)      |              | LF (u.n)     |              | HF (u.n)     |              | LF/HF        |              |
|---------------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
|                                 | RMSE         | <i>r</i>     |
| ROI_green (Li et al., 2018)▲    | 2.162        | 0.99         | 0.084        | 0.321        | 0.24         | 0.573        | 0.24         | 0.573        | 0.832        | 0.571        |
| CHROM (De Haan & Jeanne, 2013)▲ | 2.733        | 0.98         | 0.081        | 0.224        | 0.206        | 0.524        | 0.206        | 0.524        | 0.863        | 0.459        |
| POS (Wang et al., 2016)▲        | 1.906        | 0.991        | 0.07         | 0.44         | 0.158        | 0.727        | 0.158        | 0.727        | 0.679        | 0.687        |
| CVD (Niu et al., 2020)◆         | 1.26         | 0.996        | 0.058        | 0.606        | 0.09         | 0.914        | 0.09         | 0.914        | 0.453        | 0.877        |
| rPPGNet (Yu et al., 2019b)★     | 1.8          | 0.992        | 0.064        | 0.53         | 0.135        | 0.804        | 0.135        | 0.804        | 0.589        | 0.773        |
| PhysFormer (Ours)★              | <u>0.804</u> | <b>0.998</b> | <u>0.054</u> | <u>0.661</u> | <u>0.086</u> | <u>0.912</u> | <u>0.086</u> | <u>0.912</u> | <u>0.39</u>  | <u>0.896</u> |
| PhysFormer++ (Ours)★            | <b>0.765</b> | <b>0.998</b> | <b>0.052</b> | <b>0.686</b> | <b>0.083</b> | <b>0.921</b> | <b>0.083</b> | <b>0.921</b> | <b>0.368</b> | <b>0.908</b> |

The symbols ▲, ◆ and ★ denote traditional, non-end-to-end learning based and end-to-end learning based methods, respectively. Best results are marked in bold and second best in underline

**Table 6** Cross-dataset results on the MMSE-HR dataset

| Method                                | SD ↓ (bpm)  | MAE ↓ (bpm) | RMSE ↓ (bpm) | <i>r</i> ↑  |
|---------------------------------------|-------------|-------------|--------------|-------------|
| Li2014 (Li et al., 2014)▲             | 20.02       | –           | 19.95        | 0.38        |
| CHROM (De Haan & Jeanne, 2013)▲       | 14.08       | –           | 13.97        | 0.55        |
| Tulyakov2016 (Tulyakov et al., 2016)▲ | 12.24       | –           | 11.37        | 0.71        |
| ST-Attention (Niu et al., 2019b)◆     | 9.66        | –           | 10.10        | 0.64        |
| RhythmNet (Niu et al., 2019a)◆        | 6.98        | –           | 7.33         | 0.78        |
| CVD (Niu et al., 2020)◆               | 6.06        | –           | 6.04         | 0.84        |
| PhysNet (Yu et al., 2019a)★           | 12.76       | –           | 13.25        | 0.44        |
| TS-CAN (Liu et al., 2020)★            | –           | 3.85        | 7.21         | 0.86        |
| AutoHR (Yu et al., 2020)★             | 5.71        | –           | 5.87         | 0.89        |
| EfficientPhys-C (Liu et al., 2023)★   | –           | 2.91        | 5.43         | <u>0.92</u> |
| EfficientPhys-T1 (Liu et al., 2023)★  | –           | 3.48        | 7.21         | 0.86        |
| PhysFormer (Ours)★                    | <u>5.22</u> | <u>2.84</u> | <u>5.36</u>  | <u>0.92</u> |
| PhysFormer++ (Ours)★                  | <b>5.09</b> | <b>2.71</b> | <b>5.15</b>  | <b>0.93</b> |

The symbols ▲, ◆ and ★ denote traditional, non-end-to-end learning based and end-to-end learning based methods, respectively. Best results are marked in bold and second best in underline



**Fig. 7** a The scatter plot of the ground truth HR<sub>gt</sub> and the predicted HR<sub>pre</sub> via PhysFormer++ of all the face videos on VIPL-HR dataset. b Two examples of the short-time HR estimation from PhysFormer++

for face videos with significantly decreased HR. c Two example curves of the predicted rPPG signals from PhysFormer++ and the ground truth ECG signals used to calculate the HRV features

**Table 7** Ablation of Tube Tokenization of PhysFormer

| Inputs                      | [Stem] Feature size                      | [Tube size] Token numbers                      | RMSE ↓ (bpm) |
|-----------------------------|--|--|--------------|
| $160 \times 128 \times 128$ | [ $\times$ ] $160 \times 128 \times 128$ | $[4 \times 32 \times 32] 40 \times 4 \times 4$ | 10.62        |
| $160 \times 128 \times 128$ | [ $\surd$ ] $160 \times 16 \times 16$    | $[4 \times 4 \times 4] 40 \times 4 \times 4$   | <b>7.56</b>  |
| $160 \times 96 \times 96$   | [ $\surd$ ] $160 \times 12 \times 12$    | $[4 \times 4 \times 4] 40 \times 3 \times 3$   | 8.03         |
| $160 \times 128 \times 128$ | [ $\surd$ ] $160 \times 16 \times 16$    | $[4 \times 16 \times 16] 40 \times 1 \times 1$ | 10.61        |
| $160 \times 128 \times 128$ | [ $\surd$ ] $160 \times 16 \times 16$    | $[2 \times 4 \times 4] 80 \times 4 \times 4$   | 7.81         |

Best result is marked in bold

The three dimensions in tensors indicate length  $\times$  height  $\times$  width

the results of the proposed PhysFormer family and the state-of-the-art methods are shown in Table 12. It is clear that PhysFormer and PhysFormer++ generalize well in unseen domains (e.g., skin tone and lighting conditions). It is worth noting that PhysFormer++ achieves the lowest SD (5.09 bpm), MAE (2.71 bpm), RMSE (5.15 bpm) as well as the highest  $r$  (0.93) among the traditional, non-end-to-end learning and end-to-end learning based methods, indicating (1) the predicted HRs are highly correlated with the ground truth HRs, and (2) the model learns domain-invariant intrinsic rPPG-aware features. Compared with the spatio-temporal transformer based EfficientPhys-T1 (Liu et al., 2023), our proposed PhysFormer and PhysFormer++ are able to predict more accurate physiological signals, which indicates the effectiveness of the long-range spatio-temporal attention.

#### 4.5 Ablation Study

Here We provide the results of ablation studies for HR estimation on the **Fold-1** of the VIPL-HR (Niu et al., 2019a) dataset. Specifically, we first evaluate the impacts of architecture configurations for PhysFormer in terms of ‘Tube Tokenization’, ‘TD-MHSA’ and ‘ST-FF’. Then based on the optimal configuration of PhysFormer, the impacts of architecture configurations of PhysFormer++ with ‘TD-MHPSA’ and ‘SlowFast architecture’ will be studied. Finally, we study the transformer configurations (‘ $\theta$  in TDC’ and ‘layer/head numbers’) and the training receipts (‘label distribution learning’ and ‘dynamic supervision’) for the whole PhysFormer family (i.e., PhysFormer and PhysFormer++).

*Impact of Tube Tokenization in PhysFormer* In the default setting of PhysFormer, a shallow stem cascaded with a tube tokenization is used. In this ablation, we consider other four tokenization configurations with or w/o stem. It can be seen from the first row in Table 7 that the stem helps the PhysFormer see better (Xiao et al., 2021), and the RMSE increases dramatically (+ 3.06 bpm) when w/o the stem. Then we investigate the impacts of the spatial and temporal domains in tube tokenization. It is clear that the result in the fourth row with full spatial projection is quite poor (RMSE = 10.61 bpm), indicating the necessity of the spatial attention. In con-

**Table 8** Ablation of TD-MHSA and ST-FF in PhysFormer

| MHSA         | $\tau$                   | Feed-forward | RMSE (bpm) ↓ |
|--------------|--------------------------|--------------|--------------|
| –            | –                        | ST-FF        | 9.81         |
| TD-MHSA      | $\sqrt{D_h} \approx 4.9$ | ST-FF        | 9.51         |
| TD-MHSA      | 2.0                      | ST-FF        | <b>7.56</b>  |
| Vanilla MHSA | 2.0                      | ST-FF        | 10.43        |
| TD-MHSA      | 2.0                      | Vanilla FF   | 8.27         |

Best result is marked in bold

**Table 9** Ablation of TD-MHPSA for the single pathway configuration in PhysFormer++

| Pathway | MHSA     | $\mathcal{L}_{\text{atten}}$ | RMSE (bpm) ↓ |
|---------|----------|------------------------------|--------------|
| Slow    | TD-MHSA  | –                            | 7.56         |
| Slow    | TD-MHPSA | –                            | 7.69         |
| Slow    | TD-MHPSA | $\surd$                      | <b>7.43</b>  |
| Fast    | TD-MHSA  | –                            | 7.81         |
| Fast    | TD-MHPSA | –                            | 8.12         |
| Fast    | TD-MHPSA | $\surd$                      | 7.85         |

Best result is marked in bold

trast, tokenization with smaller tempos (e.g.,  $[2 \times 4 \times 4]$ ) or spatial inputs (e.g.,  $160 \times 96 \times 96$ ) reduces performance slightly. Based on the observed results, tokenizations with  $[4 \times 4 \times 4]$  and  $[2 \times 4 \times 4]$  are adopted for the defaulted setting of slow and fast pathway in PhysFormer++, respectively.

*Impact of TD-MHSA and ST-FF in PhysFormer* As shown in Table 8, both the TD-MHSA and ST-FF play vital roles in PhysFormer. The result in the first row shows that the performance degrades sharply without spatio-temporal attention. Moreover, it can be seen from the last two rows that without TD-MHSA/ST-FF, PhysFormer with vanilla MHSA/FF obtains 10.43/8.27 bpm RMSE. Thus, we can draw the conclusion that the key element ‘vanilla MHSA’ in transformer cannot provide rPPG performance gain although it captures the long-term global spatio-temporal physiological features. In contrast, the proposed ‘TD-MHSA’ benefits the rPPG measurement via local spatio-temporal physiological clue guided long-term global spatio-temporal physiological aggregation.

**Table 10** Ablation of SlowFast two-pathway based architecture in PhysFormer++

| TD-MHPSA     | Lateral connect | TD-MHCSA           | RMSE        |
|--------------|-----------------|--------------------|-------------|
| –            | –               | –                  | 7.78        |
| Slow pathway | –               | –                  | 7.58        |
| Slow pathway | High-level      | –                  | 7.34        |
| Slow pathway | Mid&High-level  | –                  | 7.38        |
| Slow pathway | High-level      | High-level         | 7.28        |
| Slow pathway | High-level      | Mid&High-level     | <b>7.16</b> |
| Slow pathway | High-level      | Low&Mid&High-level | 7.24        |

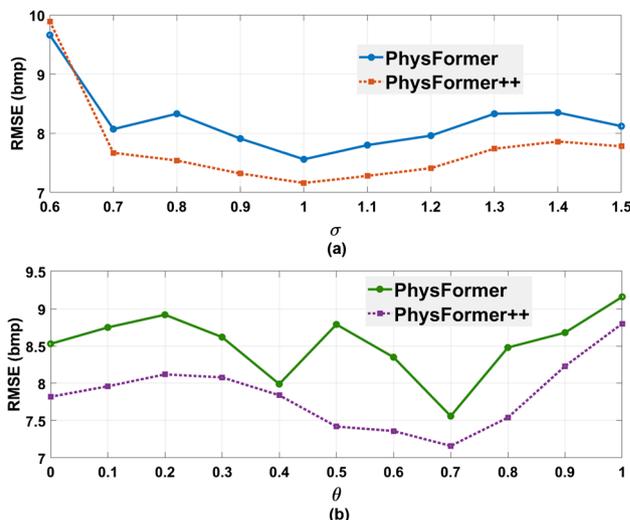
Best result is marked in bold

One important finding in this research is that, the temperature  $\tau$  influences the MHSA a lot. When the  $\tau = \sqrt{D_h}$  like previous ViT (Dosovitskiy et al., 2020; Arnab et al., 2021), the predicted rPPG signals are unsatisfied (RMSE = 9.51 bpm). Regularizing the  $\tau$  with smaller value enforces sparser spatio-temporal attention, which is effective for the quasi-periodic rPPG task.

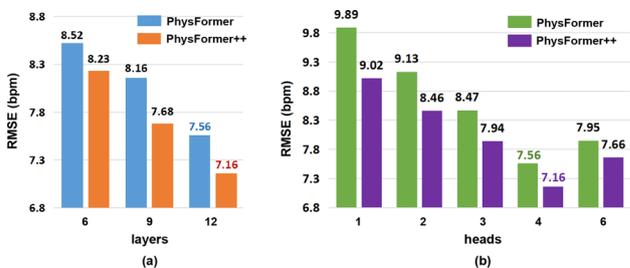
*Impact of TD-MHPSA for Different Pathway in PhysFormer++* Based on the TD-MHSA in PhysFormer, the PhysFormer++ further extends the slow pathway with the more periodic TD-MHPSA modules. Table 9 shows the results of the TD-MHPSA for single pathway configuration. It is interesting to find that compared with TD-MHSA, the performance even drop for both slow and fast pathways when assembling with TD-MHPSA but without explicit attention supervision  $\mathcal{L}_{atten}$ . When training TD-MHPSA with  $\mathcal{L}_{atten}$ , the RMSE is decreased by 0.26 and 0.27 bpm for the slow and fast pathway, respectively. It indicates the importance of explicit rPPG-aware periodicity supervision. Some visualizations with and without  $\mathcal{L}_{atten}$  can be found in Sect. 4.7.

From the results in Table 9 we can see that the TD-MHPSA with  $\mathcal{L}_{atten}$  benefits the periodic rPPG clue mining in the slow pathway while limited effects for the fast pathway. It may be because the attention loss calculated from the periodic maps with huger temporal resolution in the fast pathway is inefficient to back-propagate the rPPG-aware information. Thus, we only apply the TD-MHPSA in the slow pathway as the defaulted setting for PhysFormer++.

*Impact of the SlowFast Architecture in PhysFormer++* Table 10 illustrates the ablations of SlowFast two-pathway based architecture in PhysFormer++. From the results of the first two rows we can see that such SlowFast rPPG models even achieve inferior performance (7.78/7.58 vs. 7.56 bpm RMSE) compared with single pathway based PhysFormer. The unsatisfied results might be caused by the lack of efficient rPPG feature interaction between two pathways. We also conduct experiments with lateral connections in different levels and cross-attention based TD-MHCSA in the fast pathway. From Table 10 we can obviously find that both lateral connections and TD-MHCSA improve the performance remarkably. This is because the former one brings more temporally fine-



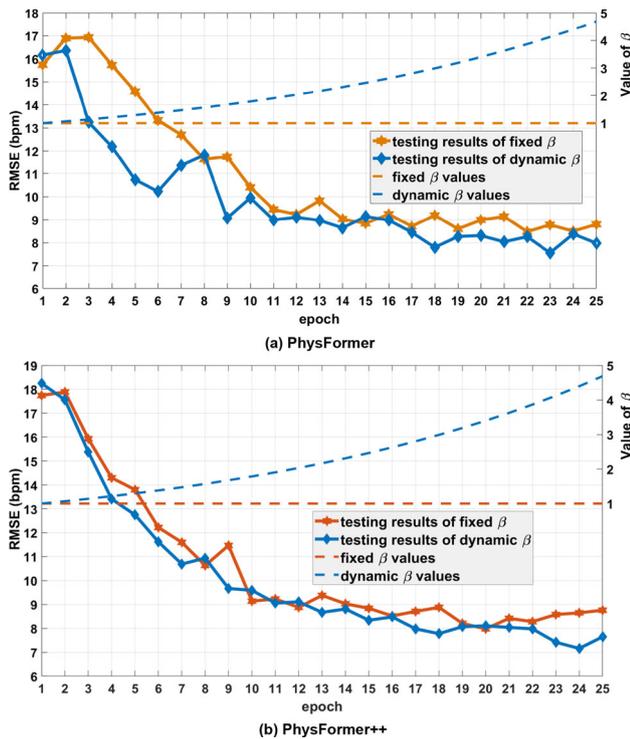
**Fig. 8** Impacts of the **a**  $\sigma$  in label distribution learning for PhysFormer and PhysFormer++ and **b**  $\theta$  in TD-MHSA, TD-MHCSA, and TD-MHPSA



**Fig. 9** Ablation of the **a** layers and **b** heads in PhysFormer and PhysFormer++

grained clues back to the slow pathway to alleviate rPPG information loss while the latter one leverages the cross-attention features to refine the redundant rPPG features in the fast pathway. The best configuration for PhysFormer++ is with high-level lateral connection and mid&high-level TD-MHCSA.

*Impact of  $\theta$  and Layer/Head Numbers in the PhysFormer Family* Hyperparameter  $\theta$  tradeoffs the contribution of local temporal gradient information. As illustrated in Fig. 8b, PhysFormer could achieve smaller RMSE when  $\theta = 0.4$



**Fig. 10** Testing results of fixed and dynamic frequency supervisions for **a** PhysFormer and **b** PhysFormer++ on the Fold-1 of VIPL-HR

and 0.7 while PhysFormer++ obtains the best performance when  $\theta = 0.7$ , indicating the importance of the normalized local temporal difference features for global spatio-temporal attention. We also investigate how the layer and head numbers influence the performance of PhysFormer and PhysFormer++. As shown in Fig. 9a, with deeper temporal transformer blocks, the RMSE are reduced progressively despite heavier computational cost. In terms of the impact of head numbers, it is clear to find from Fig. 9b that the PhysFormer family with four heads performs the best while fewer heads lead to sharp performance drops.

**Impact of Label Distribution Learning for the PhysFormer Family** Besides the temporal loss  $\mathcal{L}_{\text{time}}$  and frequency cross-entropy loss  $\mathcal{L}_{\text{CE}}$ , the ablations w/ and w/o label distribution loss  $\mathcal{L}_{\text{LD}}$  are shown in the last four rows of Table 11. Although the  $\mathcal{L}_{\text{LD}}$  performs slightly worse (respective +0.12 and +0.13 bpm RMSE for PhysFormer and PhysFormer++) than  $\mathcal{L}_{\text{CE}}$ , the best performance can be achieved using both losses, indicating the effectiveness of explicit distribution constraints for extreme-frequency interference alleviation and adjacent label knowledge propagation. It is interesting to find from the last two rows in both PhysFormer and PhysFormer++ that using real PSD distribution from ground truth PPG signals as  $\mathbf{p}$ , the performance is inferior due to the lack of an obvious peak in the distribution and partial noise.

We can also find from the Fig. 8a that the  $\sigma$  ranged from 0.9 to 1.2 for  $\mathcal{L}_{\text{LD}}$  are suitable to achieve good performance.

**Impact of Dynamic Supervision for the PhysFormer Family** Figure 10 illustrates the testing performance of PhysFormer and PhysFormer++ on Fold-1 VIPL-HR when training with fixed and dynamic supervision. It is clear that with exponentially increased frequency loss, models in the blue curves converge faster and achieve smaller RMSE. We also compare several kinds of fixed and dynamic strategies in Table 11. The results in the first four rows indicate (1) using fixed higher  $\beta$  leads to poorer performance caused by the convergency difficulty; (2) models with the exponentially increased  $\beta$  perform better than using linear increment.

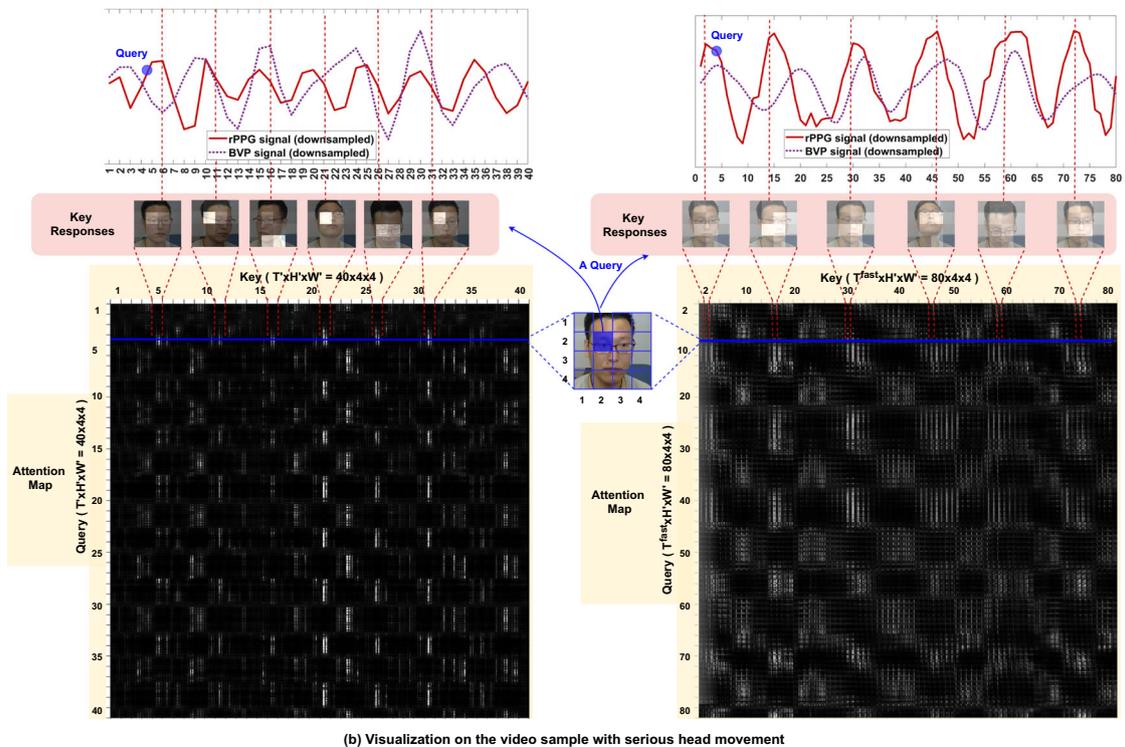
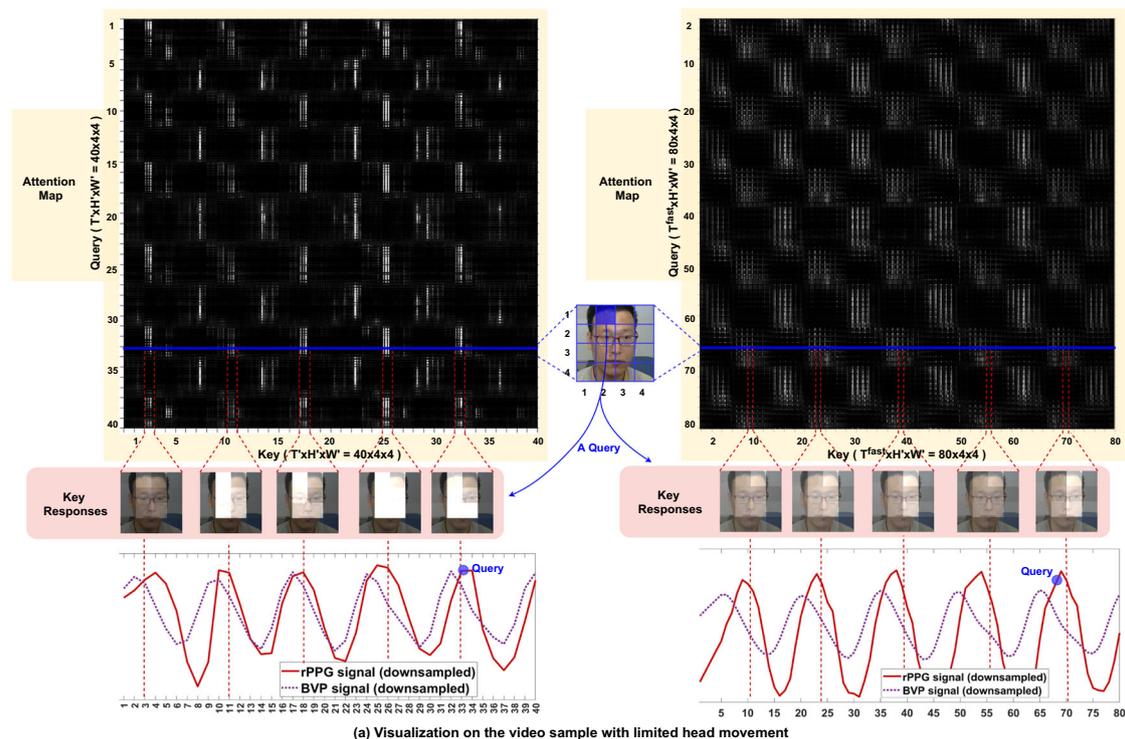
#### 4.6 Efficiency Analysis

Here we also investigate the computational cost<sup>1</sup> compared with the baselines. The number of parameters and the multiply-accumulates (MACs) are shown in Table 12. Despite huge number of parameters, PhysFormer and PhysFormer++ are with smaller MACs compared with baselines PhysNet, TS-CAN, and AutoHR. Compared with PhysFormer, the PhysFormer++ introduces extra 2.76M parameters and 1.16G MACs. The inference time of one face clip  $3 \times 160 \times 128 \times 128$  ( $C \times T \times H \times W$ ) for PhysFormer and PhysFormer++ on one V100 GPU is 29ms and 40ms, respectively. Despite slightly heavier, it can predict more accurate rPPG signals on both intra-dataset ( $-0.17$  bpm RMSE on VIPL-HR) and cross-dataset ( $-0.21$  bpm RMSE on MMSE-HR) testings. Towards efficient mobile-level rPPG applications, the computational cost of the proposed PhysFormer family is still unsatisfied. One potential future direction is to design more lightweight PhysFormer with advanced network quantization (Lin et al., 2021) and binarization (Qin et al., 2022) techniques.

#### 4.7 Visualization and Discussion

**Visualization of the Self-Attention Map** We visualize the attention maps from the last TD-MHSA module of PhysFormer (left) and the last TD-MHCSA module in the fast pathway of PhysFormer++ (right) in Fig. 11. The x and y axes of the attention map indicate the attention confidence from key and query tube tokens, respectively. From the attention maps activated from the video sample with limited head movement in Fig. 11a, we can easily find periodic or quasi-periodic responses along both axes, indicating the periodicity of the intrinsic rPPG features from PhysFormer and PhysFormer++. To be specific, given the 530th tube token (in blue) from the forehead (spatial face domain) and peak (temporal signal domain) locations as a query, the corresponding key

<sup>1</sup> <https://pypi.org/project/thop/>.



**Fig. 11** Visualization of the attention maps from (left) the 1st head in last TD-MHSA module of PhysFormer and (right) the 1st head in last TD-MHCSA module of the fast pathway in PhysFormer++. Given the 530th and 276th tube tokens in blue as the query for the video samples with **a** limited head movement and **b** serious head movement, represen-

tative key responses are illustrated (the brighter, the more attentive). The predicted downsampled rPPG signals as well as the ground truth BVP signals are shown for temporal attention understanding (Color figure online)

**Table 11** Ablation of dynamic loss in the frequency domain for PhysFormer and PhysFormer++

| Frequency loss                         | $\beta$    | Strategy    | RMSE (bpm) ↓ |
|--|------------|-------------|--------------|
| PhysFormer                             |            |             |              |
| $\mathcal{L}_{CE} + \mathcal{L}_{LD}$  | 1.0        | Fixed       | 8.48         |
| $\mathcal{L}_{CE} + \mathcal{L}_{LD}$  | 5.0        | Fixed       | 8.86         |
| $\mathcal{L}_{CE} + \mathcal{L}_{LD}$  | [1.0, 5.0] | Linear      | 8.37         |
| $\mathcal{L}_{CE} + \mathcal{L}_{LD}$  | [1.0, 5.0] | Exponential | <b>7.56</b>  |
| $\mathcal{L}_{CE}$                     | [1.0, 5.0] | Exponential | 8.09         |
| $\mathcal{L}_{LD}$                     | [1.0, 5.0] | Exponential | 8.21         |
| $\mathcal{L}_{LD}$ (real distribution) | [1.0, 5.0] | Exponential | 8.72         |
| PhysFormer++                           |            |             |              |
| $\mathcal{L}_{CE} + \mathcal{L}_{LD}$  | 1.0        | Fixed       | 7.98         |
| $\mathcal{L}_{CE} + \mathcal{L}_{LD}$  | 5.0        | Fixed       | 8.54         |
| $\mathcal{L}_{CE} + \mathcal{L}_{LD}$  | [1.0, 5.0] | Linear      | 8.13         |
| $\mathcal{L}_{CE} + \mathcal{L}_{LD}$  | [1.0, 5.0] | Exponential | <b>7.16</b>  |
| $\mathcal{L}_{CE}$                     | [1.0, 5.0] | Exponential | 7.76         |
| $\mathcal{L}_{LD}$                     | [1.0, 5.0] | Exponential | 7.89         |
| $\mathcal{L}_{LD}$ (real distribution) | [1.0, 5.0] | Exponential | 8.67         |

Best results are marked in bold

The temporal loss  $\mathcal{L}_{time}$  is with fixed  $\alpha = 0.1$  here. ‘CE’ and ‘LD’ denote cross-entropy and label distribution, respectively

**Table 12** Cross-dataset results with computational cost on MMSE-HR

| Method                             | #Param.(M) | MACs(G) | RMSE ↓ (bpm) |
|------------------------------------|------------|---------|--------------|
| PhysNet (Yu et al., 2019a)         | 0.73       | 65.19   | 13.25        |
| TS-CAN (Liu et al., 2020)          | 3.91       | 61.96   | 7.21         |
| AutoHR (Yu et al., 2020)           | 0.99       | 189.22  | 5.87         |
| EfficientPhys-C (Liu et al., 2023) | 3.84       | 31.32   | 5.43         |
| PhysFormer (Ours)                  | 7.03       | 47.01   | <u>5.36</u>  |
| PhysFormer++ (Ours)                | 9.79       | 49.85   | <b>5.15</b>  |

Best result is marked in bold and second best in underline

The FLOPs are calculated with the video input size  $3 \times 160 \times 128 \times 128$  ( $C \times T \times H \times W$ ) for PhysNet/AutoHR/PhysFormer/PhysFormer++ while  $3 \times 160 \times 96 \times 96$  for TS-CAN/EfficientPhys

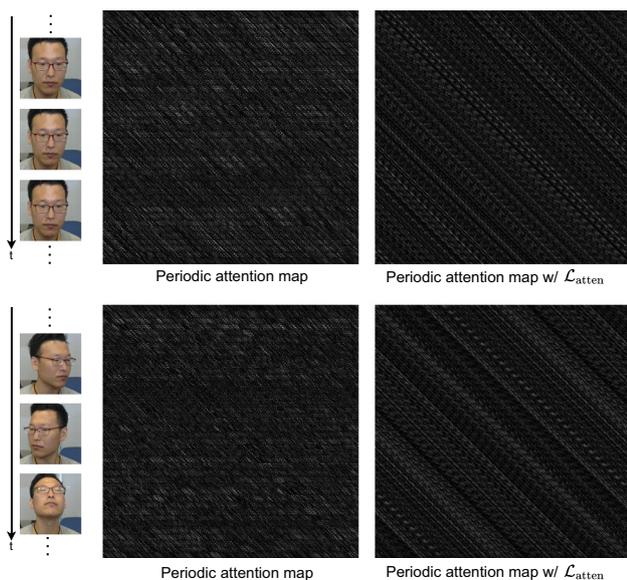
responses are illustrated at the blue line in the attention map. On the one hand, it can be seen from the key responses that dominant spatial attentions focus on the facial skin regions and discard unrelated background. On the other hand, the temporal localizations of the key responses are around peak positions in the predicted rPPG signals. All these patterns are reasonable: (1) the forehead and cheek regions (Verkruysse et al., 2008) have richer blood volume for rPPG measurement and are also reliable since these regions are less affected by facial muscle movements due to e.g., facial expressions, talking; and (2) rPPG signals from healthy people are usually periodic.

We also visualize the attention maps from another video sample with serious head movement in Fig. 11b. It can be observed from the left subfigure that the attentional response of PhysFormer is inaccurate (e.g., focusing on the neck region) when the head moves to the left. Another issue is that due to the large temporal token size ( $T_s = 4$ ) in the tok-

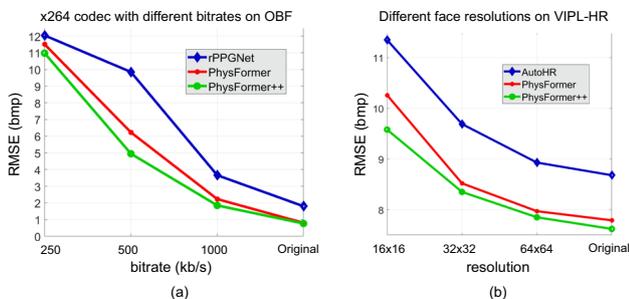
enization stage, the temporal rPPG clues might be partially discarded, resulting in the sensitivity about the head movement and the biased rPPG prediction (i.e., huge IBI gaps between the predicted rPPG and ground truth BVP signals). In contrast, it can be seen from the right subfigure in Fig. 11b that the attentional response and the predicted rPPG signal from PhysFormer++ are reliable, indicating the effectiveness of the SlowFast architecture and advanced attention modules.

Overall, two limitations of the spatio-temporal attention could be concluded from Fig. 11. First, there are still some unexpected responses (e.g., continuous query tokens with similar key responses) in the attention map, which might introduce task-irrelevant noise and damage to the performance. Second, the temporal attentions are not accurate under serious head movement scenarios, and some are coarse with phase shifts.

*Visualization of the Periodic Attention Map* We also visualize the periodic attention map from the last TD-MHPSA



**Fig. 12** Visualization of the periodic attention maps from the 1st head in last TD-MHPSA module of the slow pathway in PhysFormer++. The top row show the periodic attention map from the facial video with limited head movement while the bottom one with serious head movement



**Fig. 13** HR results with different **a** compression bitrates on OBF, and **b** resolutions on VIPL-HR

module of PhysFormer++ in Fig. 12. It is interesting to find that the periodic attention maps from the PhysFormer++ (1) trained without  $\mathcal{L}_{atten}$  are more arbitrary and easily influenced by the large head movement; and (2) trained with  $\mathcal{L}_{atten}$  are more regular and keep the periodicity even under the scenarios with serious head movement. In other words, the proposed TD-MHPSA with attention loss  $\mathcal{L}_{atten}$  enforces the PhysFormer++ to learn more periodic and robust attentional features from the face videos.

*Evaluation Under Serious Motion, Video Compression, and Low Resolution* In real-world scenarios, large head movement, high video compression rate and low face resolution usually introduce serious motion noises, compression artifacts and blurriness, respectively. All these corruptions and quality degradations make the rPPG measurement challenging. Here we evaluate the performance under these challenging scenarios.

**Table 13** HR results (RMSE (bpm)) when training with different proportion of samples on VIPL-HR

| Method                     | 10%          | 50%          | 100%        |
|----------------------------|--------------|--------------|-------------|
| AutoHR (Yu et al., 2020)   | 15.77        | <b>10.27</b> | 8.68        |
| <b>PhysFormer (Ours)</b>   | 14.84        | 11.18        | 7.79        |
| <b>PhysFormer++ (Ours)</b> | <b>13.92</b> | 10.29        | <b>7.62</b> |

Best results are marked in bold

First, we evaluate the PhysFormer family under scenarios of large head movement (i.e., ‘v2’ and ‘v9’ samples) on VIPL-HR dataset. PhysFormer and PhysFormer++ achieve RMSE of 11.46 bpm and 10.25 bpm, respectively. In other words, with richer temporally contextual rPPG clues, the two-pathway SlowFast architecture in PhysFormer++ is more motion-robust. Note that there are still performance gaps between non-end-to-end method [e.g., RhythmNet (Niu et al., 2019a) with RMSE = 9.4 bpm].

Second, we evaluate the PhysFormer family on OBF with high compression rates (250/500/1000 kb/s) using x264 codec. The corresponding HR measurement results are illustrated in Fig. 13a. Compared with the rPPGNet (Yu et al., 2019b), the PhysFormer family performs significantly better when bitrates equal to 500 and 1000 kb/s. This might be because the spatio-temporal self-attention mechanism helps filter out the compression artifacts. However, all three methods perform poorly under extremely high compression situation (i.e., bitrate = 250 kb/s).

Finally, we evaluate the PhysFormer family on VIPL-HR with different low-resolution settings to mimic the long-distance rPPG monitoring scenario. Specifically, bilinear interpolation is used to downsample the face frames to the sizes  $16 \times 16/32 \times 32/64 \times 64$  first, and then upsample them back to  $128 \times 128$ . The HR measurement results are illustrated in Fig. 13b. Despite performance drops with lower face resolution for both AutoHR (Yu et al., 2020) and the PhysFormer family, PhysFormer++ still achieves RMSE = 9.58 bpm with the lowest ( $16 \times 16$ ) resolution setting.

*Training with Fewer Samples* Since end-to-end deep models (e.g., CNNs and transformers) are data hungry, here we investigate three methods [AutoHR Yu et al. (2020), PhysFormer and PhysFormer++] under conditions of fewer training samples. As shown in Table 13, when training with only 10% or 50% samples, all these three methods obtain poor RMSE performance (> 10 bpm). Another observation is that, compared with pure CNN-based AutoHR, the proposed PhysFormer++ still achieves on par or better performance with fewer training samples. It indicates that the proposed transformer architectures can learn CNN-comparable rPPG representation even with limited data.

## 5 Conclusions

In this paper, we propose two end-to-end video transformer architectures, namely PhysFormer and PhysFormer++, for remote physiological measurement. With temporal difference transformer and elaborate supervisions, the PhysFormer family is able to achieve superior performance on benchmark datasets on both intra- and cross-testings. Comprehensive ablation studies as well as visualization analysis demonstrate the effectiveness of the proposed methods. In the future, it is potential to explore more accurate yet efficient spatio-temporal self-attention mechanism especially for long sequence rPPG monitoring. Besides the rPPG measurement task, we will investigate the effectiveness of the proposed temporal different transformer for broader fine-grained or periodic video understanding tasks in computer vision (e.g., video action recognition and repetition counting).

**Acknowledgements** This work was supported by the Academy of Finland (Academy Professor project EmotionAI with grant numbers 336116 and 345122, and ICT2023 project with grant number 345948), the National Natural Science Foundation of China (Grant No. 62002283), HKU Startup Fund, HKU Seed Fund for Basic Research, and the EPSRC grant: Turing AI Fellowship: EP/W002981/1, EPSRC/MURI grant EP/N019474/1. We would also like to thank the Royal Academy of Engineering and FiveAI. The authors wish to acknowledge CSC-IT Center for Science, Finland, for computational resources.

**Funding** Open Access funding provided by University of Oulu including Oulu University Hospital.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Arnab, A., Dehghani, M., Heigold, G., Sun, C., Lucic, M., & Schmid, C. (2021). Vivit: A video vision transformer. In *2021 IEEE/CVF international conference on computer vision (ICCV)* (pp. 6816–6826).
- Bengio, Y., Louradour, J., Collobert, R., & Weston, J. (2009). Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning* (pp. 41–48). Association for Computing Machinery.
- Bertasius, G., Wang, H., & Torresani, L. (2021). Is space–time attention all you need for video understanding? In *ICML* (Vol. 2, p. 4).
- Bulat, A., Pérez-Rúa, J.-M., Sudhakaran, S., Martínez, B., & Tzimiropoulos, G. (2021). Advances in neural information processing systems.

- In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, & J. Wortman Vaughan (Eds.), *Space-time mixing attention for video transformer* (vol. 34, pp. 19594–19607). Curran Associates, Inc.
- Cao, J., Li, Y., Zhang, K., & Gool, L. V. (2021). Video super-resolution transformer. [ArXiv:2106.06847](https://arxiv.org/abs/2106.06847)
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., & Zagoruyko, S. (2020). End-to-end object detection with transformers. [ArXiv:2005.12872](https://arxiv.org/abs/2005.12872)
- Carreira, J. & Zisserman, A. (2017). Quo vadis, action recognition? A new model and the kinetics dataset. In *2017 IEEE conference on computer vision and pattern recognition (CVPR)* (pp. 4724–4733).
- Chen, C.-F. R., Fan, Q., & Panda, R. (2021a). Crossvit: Cross-attention multi-scale vision transformer for image classification. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 357–366).
- Chen, H., Tang, H., Sebe, N., & Zhao, G. (2021b). Aniformer: Data-driven 3d animation with transformer. [ArXiv:2110.10533](https://arxiv.org/abs/2110.10533)
- Chen, H., Tang, H., Yu, Z., Sebe, N., & Zhao, G. (2021c). Geometry-contrastive transformer for generalized 3d pose transfer. In *AAAI conference on artificial intelligence*.
- Chen, W. & McDuff, D. (2018). Deepphys: Video-based physiological measurement using convolutional attention networks. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 349–365).
- Chen, X., Cheng, J., Song, R., Liu, Y., Ward, R., & Wang, Z. J. (2018). Video-based heart rate measurement: Recent advances and future prospects. *IEEE Transactions on Instrumentation and Measurement*, 68(10), 3600–3615.
- Cho, K., Van Merriënboer, B., Bahdanau, D., & Bengio, Y. (2014). On the properties of neural machine translation: Encoder–decoder approaches. [ArXiv:1409.1259](https://arxiv.org/abs/1409.1259)
- De Haan, G., & Jeanne, V. (2013). Robust pulse rate from chrominance-based rPPG. *IEEE Transactions on Biomedical Engineering*, 60(10), 2878–2886.
- Ding, M., Lian, X., Yang, L., Wang, P., Jin, X., Lu, Z., & Luo, P. (2021). Hr-nas: Searching efficient high-resolution neural architectures with lightweight transformers. In *2021 IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 2981–2991).
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., & Gelly, S., Uszkoreit, J. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. [ArXiv:2010.11929](https://arxiv.org/abs/2010.11929)
- Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., & Feichtenhofer, C. (2021). Multiscale vision transformers. In *2021 IEEE/CVF international conference on computer vision (ICCV)* (pp. 6804–6815).
- Feichtenhofer, C., Fan, H., Malik, J., & He, K. (2018). Slowfast networks for video recognition. In *2019 IEEE/CVF international conference on computer vision (ICCV)* (pp. 6201–6210).
- Gao, B.-B., Xing, C., Xie, C.-W., Wu, J., & Geng, X. (2016). Deep label distribution learning with label ambiguity. *IEEE Transactions on Image Processing*, 26, 2825–2838.
- Gao, B.-B., Zhou, H.-Y., Wu, J., & Geng, X. (2018). Age estimation using expectation of label distribution learning. In *International joint conference on artificial intelligence*.
- Geng, X., Smith-Miles, K., & Zhou, Z.-H. (2010). Facial age estimation by learning from label distributions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35, 2401–2412.
- Gideon, J. & Stent, S. (2021). 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* (pp. 3995–4004).
- Girdhar, R., Carreira, J., Doersch, C., & Zisserman, A. (2018). Video action transformer network. In *2019 IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 244–253).

- Han, K., Wang, Y., Chen, H., Chen, X., Guo, J., Liu, Z., & Yang, Z. (2022). A survey on vision transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45, 87–110.
- Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., & Wang, Y. (2021). Transformer in transformer. *Advances in Neural Information Processing Systems*, 34, 15908–15919.
- Hassani, A., Walton, S., Shah, N., Abuduweili, A., Li, J., & Shi, H. (2021). Escaping the big data paradigm with compact transformers. [ArXiv:2104.05704](https://arxiv.org/abs/2104.05704)
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *CVPR*.
- He, S., Luo, H., Wang, P., Wang, F., Li, H., & Jiang, W. (2021). Transreid: Transformer-based object re-identification. In *2021 IEEE/CVF international conference on computer vision (ICCV)* (pp. 14993–15002).
- Hsu, G.-S., Ambikapathi, A., & Chen, M.-S. (2017). Deep learning with time-frequency representation for pulse estimation from facial videos. In *2017 IEEE international joint conference on biometrics (IJCB)* (pp. 383–389). IEEE.
- Huang, C.-Z. A., Vaswani, A., Uszkoreit, J., Simon, I., Hawthorne, C., Shazeer, N. M., Dai, A. M., Hoffman, M. D., Dinculescu, M., & Eck, D. (2019). Music transformer: Generating music with long-term structure. In *International conference on learning representations*.
- Kazakos, E., Nagrani, A., Zisserman, A., & Damen, D. (2021). Slow-fast auditory streams for audio recognition. In *ICASSP 2021—2021 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 855–859).
- Khan, S., Naseer, M., Hayat, M., Zamir, S. W., Khan, F. S., & Shah, M. (2021). Transformers in vision: A survey. [ArXiv:2101.01169](https://arxiv.org/abs/2101.01169)
- Lam, A. & Kuno, Y. (2015). Robust heart rate measurement from video using select random patches. In *Proceedings of the IEEE international conference on computer vision* (pp. 3640–3648).
- Lee, E., Chen, E., & Lee, C.-Y. (2020). Meta-rppg: Remote heart rate estimation using a transductive meta-learner. In *European conference on computer vision*.
- Li, X., Alikhani, I., Shi, J., Seppänen, T., Junttila, J. M., Majamaa-Voltti, K., Tulppo, M. P., & Zhao, G. (2018). The obf database: A large face video database for remote physiological signal measurement and atrial fibrillation detection. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)* (pp. 242–249).
- Li, X., Chen, J., Zhao, G., & Pietikainen, M. (2014). Remote heart rate measurement from face videos under realistic situations. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 4264–4271).
- Lin, J., Gan, C., & Han, S. (2018). Tsm: Temporal shift module for efficient video understanding. In *2019 IEEE/CVF international conference on computer vision (ICCV)* (pp. 7082–7092).
- Lin, T., Wang, Y., Liu, X., & Qiu, X. (2022). *A survey of transformers*. AI Open.
- Lin, Y., Zhang, T., Sun, P., Li, Z., & Zhou, S. (2021). Fq-vit: Fully quantized vision transformer without retraining. [ArXiv:2111.13824](https://arxiv.org/abs/2111.13824)
- Liu, R., Deng, H., Huang, Y., Shi, X., Lu, L., Sun, W., Wang, X., Dai, J., & Li, H. (2021a). Fuseformer: Fusing fine-grained information in transformers for video inpainting. In *2021 IEEE/CVF international conference on computer vision (ICCV)* (pp. 14020–14029).
- Liu, X., Fromm, J., Patel, S., & McDuff, D. (2020). Multi-task temporal shift attention networks for on-device contactless vitals measurement. *Advances in Neural Information Processing Systems*, 33, 19400–19411.
- Liu, X., Hill, B., Jiang, Z., Patel, S., & McDuff, D. (2023). Efficient-phys: Enabling simple, fast and accurate camera-based cardiac measurement. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision* (pp. 5008–5017).
- Liu, X., Patel, S., & McDuff, D. (2021b). Camera-based physiological sensing: Challenges and future directions. [ArXiv:2110.13362](https://arxiv.org/abs/2110.13362)
- Liu, X., Wang, Q., Hu, Y., Tang, X., Zhang, S., Bai, S., & Bai, X. (2021c). End-to-end temporal action detection with transformer. *IEEE Transactions on Image Processing*, 31, 5427–5441.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., & Guo, B. (2021d). Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 10012–10022).
- Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., & Hu, H. (2021e). Video swin transformer. In *2022 IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 3192–3201).
- Lu, H., & Han, H. (2021). Nas-hr: Neural architecture search for heart rate estimation from face videos. *Virtual Reality & Intelligent Hardware*, 3(1), 33–42.
- Lu, H., Han, H., & Zhou, S. K. (2021). Dual-gan: Joint bvp and noise modeling for remote physiological measurement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 12404–12413).
- Magdalena Nowara, E., Marks, T. K., Mansour, H., & Veeraraghavan, A. (2018). Sparseppg: Towards driver monitoring using camera-based vital signs estimation in near-infrared. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops* (pp. 1272–1281).
- Neimark, D., Bar, O., Zohar, M., & Asselmann, D. (2021). Video transformer network. In *2021 IEEE/CVF international conference on computer vision workshops (ICCVW)* (pp. 3156–3165).
- Niu, X., Han, H., Shan, S., & Chen, X. (2017). Continuous heart rate measurement from face: A robust rppg approach with distribution learning. In *2017 IEEE international joint conference on biometrics (IJCB)* (pp. 642–650).
- Niu, X., Han, H., Shan, S., & Chen, X. (2018). Synrhythm: Learning a deep heart rate estimator from general to specific. In *2018 24th international conference on pattern recognition (ICPR)* (pp. 3580–3585). IEEE.
- Niu, X., Shan, S., Han, H., & Chen, X. (2019a). Rhythmnet: End-to-end heart rate estimation from face via spatial-temporal representation. *IEEE Transactions on Image Processing*, 29, 2409–2423.
- Niu, X., Yu, Z., Han, H., Li, X., Shan, S., & Zhao, G. (2020). Video-based remote physiological measurement via cross-verified feature disentangling. In *ECCV* (pp. 295–310). Springer.
- Niu, X., Zhao, X., Han, H., Das, A., Dantcheva, A., Shan, S., & Chen, X. (2019b). Robust remote heart rate estimation from face utilizing spatial-temporal attention. In *2019 14th IEEE international conference on automatic face & gesture recognition (FG 2019)* (pp. 1–8). IEEE.
- Nowara, E. M., McDuff, D., & Veeraraghavan, A. (2021). The benefit of distraction: Denoising camera-based physiological measurements using inverse attention. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 4955–4964).
- Poh, M.-Z., McDuff, D. J., & Picard, R. W. (2010a). Advancements in noncontact, multiparameter physiological measurements using a webcam. *IEEE Transactions on Biomedical Engineering*, 58(1), 7–11.
- Poh, M.-Z., McDuff, D. J., & Picard, R. W. (2010b). Non-contact, automated cardiac pulse measurements using video imaging and blind source separation. *Optics Express*, 18(10), 10762–10774.
- Qin, H., Ding, Y., Zhang, M., Yan, Q., Liu, A., Dang, Q., Liu, Z., & Liu, X. (2022). Bibert: Accurate fully binarized bert. In *ICLR*.
- Qiu, Y., Liu, Y., Artega-Falconi, J., Dong, H., & El Saddik, A. (2018). EVM-CNN: Real-time contactless heart rate estimation from facial video. *IEEE Transactions on Multimedia*, 21(7), 1778–1787.
- Revanur, A., Dasari, A., Tucker, C. S., & Jeni, L. A. (2022). Instantaneous physiological estimation using video transformers. [ArXiv:2202.12368](https://arxiv.org/abs/2202.12368)

- Shaw, P., Uszkoreit, J., & Vaswani, A. (2018). Self-attention with relative position representations. In *North American chapter of the Association for Computational Linguistics*.
- Soleymani, M., Lichtenauer, J., Pun, T., & Pantic, M. (2012). A multimodal database for affect recognition and implicit tagging. *IEEE Transactions on Affective Computing*, 3, 42–55.
- Špetlík, R., Franc, V., & Matas, J. (2018). Visual heart rate estimation with convolutional neural network. In *Proceedings of the British machine vision conference*, Newcastle, UK (pp. 3–6).
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., & Jégou, H. (2021). Training data-efficient image transformers & distillation through attention. In *International conference on machine learning* (pp. 10347–10357). PMLR.
- Tulyakov, S., Alameda-Pineda, X., Ricci, E., Yin, L., Cohn, J. F., & Sebe, N. (2016). Self-adaptive matrix completion for heart rate estimation from face videos under realistic conditions. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2396–2404).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, 30.
- Verkruysse, W., Svaasand, L. O., & Nelson, J. S. (2008). Remote plethysmographic imaging using ambient light. *Optics Express*, 16(26), 21434–21445.
- Wang, L., Yang, H., Wu, W., Yao, H., & Huang, H. (2021a). Temporal action proposal generation with transformers. [ArXiv:2105.12043](https://arxiv.org/abs/2105.12043)
- Wang, W., Den Brinker, A. C., Stuijk, S., & De Haan, G. (2016). Algorithmic principles of remote PPG. *IEEE Transactions on Biomedical Engineering*, 64(7), 1479–1491.
- Wang, W., Xie, E., Li, X., Fan, D.-P., Song, K., Liang, D., Lu, T., Luo, P., & Shao, L. (2021b). Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 568–578).
- Wu, K., Peng, H., Chen, M., Fu, J., & Chao, H. (2021). Rethinking and improving relative position encoding for vision transformer. In *2021 IEEE/CVF international conference on computer vision (ICCV)* (pp. 10013–10021).
- Xiao, T., Singh, M., Mintun, E., Darrell, T., Dollár, P., & Girshick, R. B. (2021). Early convolutions help transformers see better. In *Neural information processing systems*.
- Xu, M., Xiong, Y., Chen, H., Li, X., Xia, W., Tu, Z., & Soatto, S. (2021). Long short-term transformer for online action detection. [ArXiv:2107.03377](https://arxiv.org/abs/2107.03377)
- Yu, Z., Li, X., Niu, X., Shi, J., & Zhao, G. (2020). Autohr: A strong end-to-end baseline for remote heart rate measurement with neural searching. *IEEE Signal Processing Letters*, 27, 1245–1249.
- Yu, Z., Li, X., Wang, P., & Zhao, G. (2021). Transrppg: Remote photoplethysmography transformer for 3d mask face presentation attack detection. *IEEE Signal Processing Letters*, 28, 1290–1294.
- Yu, Z., Li, X., & Zhao, G. (2019a). Remote photoplethysmograph signal measurement from facial videos using spatio-temporal networks. In *British machine vision conference* (pp. 277–289).
- Yu, Z., Li, X., & Zhao, G. (2021). Facial-video-based physiological signal measurement: Recent advances and affective applications. *IEEE Signal Processing Magazine*, 38(6), 50–58.
- Yu, Z., Peng, W., Li, X., Hong, X., & Zhao, G. (2019b). 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* (pp. 151–160).
- Yu, Z., Qin, Y., Li, X., Zhao, C., Lei, Z., & Zhao, G. (2021). Deep learning for face anti-spoofing: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Yu, Z., Shen, Y., Shi, J., Zhao, H., Torr, P. H., & Zhao, G. (2022). Physformer: facial video-based physiological measurement with temporal difference transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 4186–4196).
- Yu, Z., Zhou, B., Wan, J., Wang, P., Chen, H., Liu, X., Li, S. Z., & Zhao, G. (2022). Deep learning for face anti-spoofing: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. <https://doi.org/10.1109/TPAMI.2022.3215850>.
- Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z.-H., Tay, F. E., Feng, J., & Yan, S. (2021). Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 558–567).
- Zeng, Y., Fu, J., & Chao, H. (2020). Learning joint spatial-temporal transformations for video inpainting. [ArXiv:2007.10247](https://arxiv.org/abs/2007.10247)
- Zhang, K., Zhang, Z., Li, Z., & Qiao, Y. (2016). Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE SPL*, 23, 1499–1503.
- Zhao, J., Li, X., Liu, C., Shuai, B., Chen, H., Snoek, C. G. M., & Tighe, J. (2021). Tuber: Tube-transformer for action detection. [ArXiv:2104.00969](https://arxiv.org/abs/2104.00969)
- Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P. H. S., & Zhang, L. (2020). Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *2021 IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 6877–6886).

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.