

Joint Spatial-Temporal Modeling and Contrastive Learning for Self-supervised Heart Rate Measurement

Wei Qian^{1,†}, Qi Li^{3,4,†}, Kun Li^{5,*}, Xinke Wang^{4,3}, Xiao Sun^{1,2,3}, Meng Wang^{1,2,3} and Dan Guo^{1,2,3,6,*}

¹School of Computer Science and Information Engineering, School of Artificial Intelligence, Hefei University of Technology (HFUT)

²Key Laboratory of Knowledge Engineering with Big Data (HFUT), Ministry of Education

³Institute of Artificial Intelligence, Hefei Comprehensive National Science Center, China

⁴Anhui University, China

⁵Zhejiang University, China

⁶Anhui Zhonghuitong Technology Co., Ltd.

Abstract

This paper briefly introduces the solutions developed by our team, HFUT-VUT, for Track 1 of self-supervised heart rate measurement in the 3rd Vision-based Remote Physiological Signal Sensing (RePSS) Challenge hosted at IJCAI 2024. The goal is to develop a self-supervised learning algorithm for heart rate (HR) estimation using unlabeled facial videos. To tackle this task, we present two self-supervised HR estimation solutions that integrate spatial-temporal modeling and contrastive learning, respectively. Specifically, we first propose a non-end-to-end self-supervised HR measurement framework based on spatial-temporal modeling, which can effectively capture subtle rPPG clues and leverage the inherent bandwidth and periodicity characteristics of rPPG to constrain the model. Meanwhile, we employ an excellent end-to-end solution based on contrastive learning, aiming to generalize across different scenarios from complementary perspectives. Finally, we combine the strengths of the above solutions through an ensemble strategy to generate the final predictions, leading to a more accurate HR estimation. As a result, our solutions achieved a remarkable RMSE score of 8.85277 on the test dataset, securing **2nd place** in Track 1 of the challenge.

Keywords

Self-supervised, heart rate, rPPG, spatial-temporal modeling, contrastive learning

1. Introduction

Remote physiological measurement [1, 2, 3, 4, 5] has emerged as a promising field with significant applications in healthcare, wellness monitoring, and human-computer interaction.


The 3rd Vision-based Remote Physiological Signal Sensing (RePSS) Challenge & Workshop, Aug 3–9, 2024, Jeju, South Korea

*Corresponding authors.

†These authors contributed equally.

✉ qianwei.hfut@gmail.com (W. Qian); liqi@stu.ahu.edu.cn (Q. Li); kunli.hfut@gmail.com (K. Li); xinkewang689@gmail.com (X. Wang); sunx@hfut.edu.cn (X. Sun); eric.mengwang@gmail.com (M. Wang); guodan@hfut.edu.cn (D. Guo)

ORCID: 0009-0007-9467-6296 (W. Qian); 0000-0002-8655-5781 (Q. Li); 0000-0001-5083-2145 (K. Li); 0009-0002-8399-8322 (X. Wang); 0000-0001-9750-7032 (X. Sun); 0000-0002-3094-7735 (M. Wang); 0000-0003-2594-254X (D. Guo)

 © 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Traditional methods for physiological measurement, such as electrocardiograms (ECG) and photoplethysmograms (PPG), require direct contact with the skin, which can be cumbersome and inconvenient for continuous monitoring. With the great success of deep learning in computer vision [6, 7, 8, 9, 10], recent advancements [11, 12] have paved the way for non-contact, video-based techniques to estimate physiological signals such as heart rate (HR) and respiratory rate (RR) from facial videos, providing a more comfortable and accessible approach for users.

Despite the promising potential of video-based physiological measurement, most existing methods [13, 5, 3] rely heavily on supervised learning, necessitating large amounts of labeled data for training. Acquiring such labeled data is often labor-intensive and time-consuming, posing a significant bottleneck for developing robust and generalizable models. Moreover, supervised methods may not generalize well across different environments and lighting conditions, limiting their practical applicability. Therefore, the development of label-free rPPG estimation methods is becoming increasingly urgent.

To address these challenges, the 3rd Vision-based Remote Physiological Signal Sensing (RePSS) Challenge at IJCAI 2024 was launched. This challenge aims to develop self-supervised training methods for HR measurement using unlabeled facial videos, thereby reducing the dependency on extensive labeled datasets. For this challenge, we present two self-supervised HR estimation solutions that integrate spatial-temporal modeling and contrastive learning, respectively. Inspired by Dual-TL [3] and SiNC [14], we propose a non-end-to-end self-supervised HR measurement framework based on a spatial-temporal Transformer to capture subtle rPPG clues. Meanwhile, we adopt a complementary end-to-end contrastive learning solution based on Contrast-Phys+ [11] to enhance the model accuracy. Finally, we combine the strengths of both solutions through an ensemble strategy to generate the final predictions, securing second place with the RMSE score of 8.85277.

In conclusion, the main contributions can be summarized as follows:

- We propose a non-end-to-end self-supervised solution based on spatial-temporal modeling. By considering the priors of periodicity consistency and bandwidth limitation of the rPPG signal, we introduce four loss functions to supervise the model effectively.
- We present an end-to-end solution based on contrastive learning, which utilizes 3DCNN to extract features and employs a contrastive loss to learn discriminative representations for periodic rPPG signal modeling.
- Our solution achieved second place with the RMSE score of 8.85277 on the test dataset in Track 1 of the 3rd Vision-based Remote Physiological Signal Sensing Challenge. The experimental results demonstrate the effectiveness and robustness of our proposed solutions.

2. Methodology

2.1. Solution 1: Self-supervised HR Measurement with Spatial-Temporal Transformer

Inspired by the great success of Transformer in computer vision [15], we present a non-end-to-end self-supervised HR measurement framework to mitigate the need for labeled video




Figure 1: Overview of the proposed solution 1. Given an input facial video with T frames, we obtain N facial ROIs for each frame and extract the MSTmap representation $M \in \mathbb{R}^{T \times N \times C}$ for the video, where N is the number of facial ROI. A feature embedding layer is used to project the MSTmap to high-dimensional feature $X \in \mathbb{R}^{T \times N \times D}$. Then, we stack spatial-temporal Transformer for L loops to capture subtle rPPG clues. Next, a rPPG regression head is used to output rPPG signal $s_{pre} \in \mathbb{R}^{T \times 1}$. Finally, we apply four self-supervised losses to constrain the model.

data based on a Spatial-Temporal Transformer. The overview of this solution is illustrated in Figure 1. Specifically, we first transform the input facial video into a multi-scale spatial-temporal map (MSTmap) in Section 2.1.1. Then, we introduce our spatial-temporal Transformer module in Section 2.1.2. Next, in Section 2.1.3, with the constraints of periodicity consistency and bandwidth finiteness, our model directly discovers blood volume pulses from unlabeled videos to predict HR.

2.1.1. Data Pre-processing

The quasi-periodic pulse signal originates from subtle light reflections of blood vessels under the skin. Therefore, non-skin pixels and facial geometric features can be considered as rPPG-independent noises. We transform the raw facial video into MSTmap to highlight the spatiotemporal information of the human face, which is a common practice in rPPG measurement [16, 17]. Concretely, the MSTmap divides the facial area into 6 meta-ROI blocks, which can generate $N = (2^6 - 1) = 63$ ROI combination blocks, and the pixels of each block are averaged separately for C color channels. In the video, all the frames are concatenated along the time dimension to generate a spatial-temporal map of size $\mathbb{R}^{T \times N \times C}$, where $C = 6$ represents {R,G,B,Y,U,V} channels. Next, we embed the MSTmap M to high-dimensional feature $X \in \mathbb{R}^{T \times N \times D}$ with feature dimension D by using a full-connected layer.

2.1.2. Spatial-Temporal Transformer

Our spatial-temporal Transformer tailored for remote physiological measurement is designed carefully for perceiving the temporal and spatial correlations. It includes two encoders (spatial encoder and temporal encoder) to refine the ROI representation containing rPPG clues by capturing long-term spatiotemporal contextual information. We now explain the proposed model in detail. Specifically, given the input features $X \in \mathbb{R}^{T \times N \times D}$, the process of embedding

spatial context for t -frame can be formulated as:

$$\begin{aligned} \mathbf{Q}^{(t)} &= \mathbf{X}^{(t)}W_{tq}, \mathbf{K}^{(t)} = \mathbf{X}^{(t)}W_{tk}, \mathbf{V}^{(t)} = \mathbf{X}^{(t)}W_{tv}, \\ \mathbf{Z}^{(t)} &= \text{softmax}\left(\frac{\mathbf{Q}^{(t)}\mathbf{K}^{(t)T}}{\sqrt{D}}\right)\mathbf{V}^{(t)} + \mathbf{X}^{(t)}, \\ \mathbf{Z}'^{(t)} &= \text{MLP}(\text{LN}(\mathbf{Z}^{(t)})) + \mathbf{Z}^{(t)}, \end{aligned} \quad (1)$$

where W_{tq}, W_{tk}, W_{tv} are learnable parameters shaped as $D \times D$. $\mathbf{X}^{(t)}$ denote the feature in t -th frame. MLP is the multi-layer perceptron layer and LN is layer normalization operation. The feature map of all frames $\{\mathbf{Z}^{(t)}|t = 1, \dots, T\}$ are concatenated together into $\mathbf{Z}_s \in \mathbb{R}^{T \times N \times D}$.

The other complementary module is applied to enhance the input rPPG features with temporal dynamical transition clues and enrich the temporal context by highlighting the informative features along the time dimension for each facial ROI. Our temporal encoder follows the way in Eq. 1. The difference is that we calculate the temporal dimension for each spatial unit ($n \in [1, N]$). We output the temporally correlated feature for the n -th facial ROI feature as $\mathbf{Z}'^{(n)} \in \mathbb{R}^{T \times D}$ and stack the features $\{\mathbf{Z}'^{(n)}|n = 1, 2, \dots, N\}$ together, represented by $\mathbf{Z}_t \in \mathbb{R}^{N \times T \times D}$.

The spatial and temporal encoders are stacked as L loops in an alternating manner, taking into account the spatial and temporal complementary contextual information integrally. Moreover, spatial and temporal position embedding is applied only to the first encoder to retain two kinds of position information. Finally, we use an rPPG regression head to project the feature to a 1D rPPG signal $y_{pred} \in \mathbb{R}^{T \times 1}$.

2.1.3. Self-supervised Loss

As highlighted in previous studies [18, 14], the rPPG signal possesses inherent theoretical priors, including specific bandwidth in the frequency domain. By incorporating this prior knowledge, we employ three self-supervised loss functions from [14] in this work. Additionally, to further effectively train the model, we also propose a new periodicity loss based on periodic characteristics of the rPPG signal. Notably, all predicted rPPG signals are transformed into power spectrum density (PSD) with the Fast Fourier Transform (FFT) before computing all losses in our method, denoted as $F = \text{FFT}(y)$.

Bandwidth Loss. A healthy HR falls within a specific frequency range. Following the [14], we penalize the model for producing signals that exceed the healthy HR bandwidth limits. Consequently, the bandwidth loss can be formalized as follows:

$$\mathcal{L}_{band} = \frac{1}{\sum_{i=-\infty}^{\infty} F_i} \left[\sum_{i=-\infty}^a F_i + \sum_{i=b}^{\infty} F_i \right], \quad (2)$$

where a and b denote lower and upper band limits, respectively. F_i is the power in the i th frequency bin of the predicted signal. In our experiments, we specify the limits as $a = 0.66$ Hz to $b = 3$ Hz, which corresponds to a common pulse rate range from 40 bpm to 180 bpm. This range effectively captures the typical variations in a healthy HR, ensuring that our model focuses on the relevant frequency components while minimizing the influence of noise. By incorporating

this bandwidth loss, our model is better equipped to distinguish between meaningful rPPG signals and disturbances, ultimately leading to more accurate HR estimation.

Sparsity Loss. Since we are primarily interested in heartbeat frequency, we emphasize the periodic heartbeats by suppressing non-heartbeat frequencies. Following [14], we penalize the energy in the bandwidth regions far away from the spectrum peak, which can ensure that the model focuses on the relevant heartbeat frequencies. It can be formulated as:

$$\mathcal{L}_{sparse} = \frac{1}{\sum_{i=a}^b F_i} \left[\sum_{i=a}^{\text{argmax}(F) - \Delta_F} F_i + \sum_{i=\text{argmax}(F) + \Delta_F}^b F_i \right], \quad (3)$$

where $\text{argmax}(F)$ is the frequency of the spectral peak, and $\Delta_F = 6$ is the frequency padding around the peak. This loss enhances the model's ability to accurately estimate HR by ensuring that the spectral energy is concentrated around the true HR frequencies, thus minimizing the influence of noise and other non-relevant frequency components.

Variance Loss. To avoid the model collapsing to a specific frequency, we also use a variance loss [14, 19] to spread the variance of the power spectral density into a uniform distribution over the desired frequency band. Firstly, we define a uniform prior distribution P over d frequencies. Then, we consider a batch of n spectral densities, represented as $F = [v_1, \dots, v_n]$, where each v_i is a d -dimensional frequency decomposition of a predicted waveform. To aggregate these spectral densities, we compute the normalized sum across the batch, denoted as Q . Therefore, the variance loss \mathcal{L}_{var} can be formulated as:

$$\mathcal{L}_{var} = \frac{1}{d} \sum_{i=1}^d (\text{CDF}_i(Q) - \text{CDF}_i(P))^2, \quad (4)$$

where CDF_i represents the cumulative distribution function at the i -th frequency.

Periodicity Loss. In addition to the intrinsic properties of the rPPG signal itself, we have observed that adjacent rPPG signals do not change rapidly over short periods. This is typically manifested by similar periodicity in neighboring rPPG signals, meaning they share a dominant peak in the PSD. Specifically, we uniformly sample S non-overlapping temporal segments from a short rPPG signal (e.g., 10s). The PSDs of these segments should be similar. Thus, our proposed periodicity loss can be formulated as:

$$\mathcal{L}_{perio} = \sum_{j=1}^{S-1} \sum_{i=-\infty}^{\infty} (F_i^j - F_i^{j+1})^2, \quad (5)$$

where $S = 3$ denotes the number of segments.

In summary, the overall loss function of our self-supervised learning strategy is :

$$\mathcal{L}_{total} = \mathcal{L}_{band} + \mathcal{L}_{sparse} + \mathcal{L}_{var} + \mathcal{L}_{perio}. \quad (6)$$




Figure 2: Overview of the solution 2. In the pre-train stage, the model is trained in a contrastive learning-based self-supervised manner. After that, the pre-trained model is fine-tuned by supervised loss.

2.2. Solution 2: Self-supervised HR Measurement with Contrastive Learning

Here we provide the end-to-end self-supervised HR measurement framework based on the contrastive learning strategy. The framework is depicted in Figure 2. Specifically, we first perform data-preprocessing in Section 2.2.1. Then we pre-train the proposed model in an unsupervised setting based on the Contrast-Phys+ [11] in Section 2.2.2. Finally, we fine-tune the Contrast-Phys+ model with a supervised setting and obtain the final rPPG predictor in Section 2.2.3.

2.2.1. Data Pre-processing

In this self-supervised manner, we input facial video into our model to estimate the final rPPG signal. For an original video, we first perform face detection by MTCNN [20] to get the four coordinates of the face bounding box from the first frame. Then, we enlarge the length and width of the bounding box by 1.5 times and crop the face region for each frame of the video. The cropped faces are resized to 128×128 . Next, we segment each video into clips to feed into the model. Note that we also perform frame difference operations on the clip to generate normalized difference frames as an attempt of model input. The difference between two consecutive frames can be formulated as:

$$\Delta V_t = V_{t+1} - V_t, \quad (7)$$

where V_t denotes the t -th frame. To keep the length of the difference video equal to the raw video, we simply repeat the last difference frame. Then, the ΔV is normalized.

2.2.2. Pre-training

In this stage, following the setting of [11] we modify the 3DCNN-based PhysNet to get spatiotemporal rPPG (ST-rPPG) block representation. The model outputs spatiotemporal rPPG features with shape $T \times S \times S$, where T is the temporal length, and S is the spatial dimension. The ST-rPPG block can be regarded as a collection of rPPG signals from different facial regions. Therefore, for each input, we can sample S^2 rPPG signals with the length of T .

According to the observations that rPPG spatial similarity and temporal similarity in [11], the ST-rPPG block can sample multiple rPPG signals with short time intervals and different spatial

positions. Those signals should be similar. Then contrastive learning can be formulated by pulling together the rPPG signals from the same ST-rPPG block and pushing away the signals from different ST-rPPG blocks extracted in the crossing video. The contrastive loss can be formulated as:

$$\mathcal{L}_{pos} = \sum_{i=1}^N \sum_{\substack{j=1 \\ j \neq i}}^N \left(\|f_i - f_j\|^2 + \|f'_i - f'_j\|^2 \right) / (2N(N-1)), \quad (8)$$

$$\mathcal{L}_{neg} = - \sum_{i=1}^N \sum_{j=1}^N \|f_i - f'_j\|^2 / N^2, \quad (9)$$

$$\mathcal{L}_{ctr} = \mathcal{L}_{pos} + \mathcal{L}_{neg}, \quad (10)$$

where f_i denotes the Power Spectrum Densities (PSDs) of the rPPG signal in position i and f'_i is the other video's PSDs. N is the number of sampled rPPG pairs. The contrastive loss function minimizes the MSE distance between positive samples and maximizes the distance between the negative samples to force the model to learn the discriminative representation of the underlying signals from different videos.

2.2.3. Fine-tuning

With the pre-trained 3DCNN-based PhysNet model, we use the officially designated dataset to fine-tune it in a supervised manner. Specifically, in this stage, we modified the output of the model by averaging the spatial dimension and then obtained a predicted rPPG signal. Given the predicted rPPG signal y_{pred} and the ground-truth PPG signal y_{gt} , a popular Negative Pearson correlation (Pear) loss and Negative max cross-correlation (MCC) loss are selected to perform supervised training. It is worth noting that the Pear is the time domain loss function while the MCC loss is the frequency domain loss function. The MCC is robust to temporal offsets in the ground truth, which can make up for the Pear loss. The MCC loss is formulated as:

$$\mathcal{L}_{mcc} = - \text{Max} \left(\frac{FFT^{-1}\{BPass(FFT\{y_{pred}\} \cdot \overline{FFT\{y_{gt}\}})\}}{\sigma_{y_{pred}} \times \sigma_{y_{gt}}} \right), \quad (11)$$

where FFT^{-1} is the inverse of fast Fourier transform (FFT), σ is the standard deviation. Besides, as the ground-truth signals are the reference of predicted rPPG signals, the y_{pred} should be similar to y_{gt} . Therefore, we also use the contrastive loss by the following:

$$\mathcal{L}_{pos}^{gt} = \sum_{i=1}^N \sum_{\substack{j=1 \\ j \neq i}}^N \left(\|f_i - g_j\|^2 + \|f'_i - g'_j\|^2 \right) / (2N(N-1)), \quad (12)$$

$$\mathcal{L}_{neg}^{gt} = - \sum_{i=1}^N \sum_{j=1}^N \left(\|f_i - g'_j\|^2 + \|f'_i - g_j\|^2 \right) / N^2, \quad (13)$$

where g is the PSDs of the ground-truth signal.

Finally, the overall loss for fine-tuning is the combination of Pear loss, MCC loss, and contrastive loss, which can resist noise interference of ground-truth signal.

$$\mathcal{L}_s = \mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt} + \alpha \mathcal{L}_{pear} + \beta \mathcal{L}_{mcc}, \quad (14)$$

where \mathcal{L}_{pear} is the Negative Pearson correlation loss function. In our experiments, we set α to 0.1 and β to 0.2 for the VIPL-V2 dataset.

3. Experiments

3.1. Datasets

UBFC-rPPG [21] is a commonly used pure dataset for physiological estimation. It records 42 facial videos from 42 subjects in a stable lab environment. **PURE** [22] contains 60 facial videos of 10 participants under 6 modes (steady, small rotation, medium rotation, talking, slow translation, and fast translation). **MMSE-HR** [23] contains 102 facial videos captured from 40 subjects under six task modes. This dataset contains various facial expression changes. **DISFA** [24] is a non-posed facial expression dataset. It records 27 facial videos from 27 subjects with different ethnicities[25]. **VIPL-V2** [26] is the second version of the VIPL-HR [26] dataset for remote HR estimation from face videos under less-constrained situations, which contains 2,000 RGB videos provided in this challenge [16, 17]. Up until the publication of the **OBF** [2] dataset, it contains 100 healthy subjects and 6 patients with atrial fibrillation, totaling 10,600 minutes in length [13]. In this challenge, some data of OBF are included in the test set. Following the rule of this challenge, we use the datasets except VIPL-V2 and OBF without labels to pre-train the model and finetune the model on the VIPL-V2 dataset.

3.2. Evaluation Metrics and Implementation Details

In this challenge, the root mean squared error (RMSE) is selected as the evaluation metric between the predicted HR y_{pred} and ground-truth HR y_{gt} as below:

$$RMSE(y_{pred}, y_{gt}) = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_{pred}^i - y_{gt}^i)^2}, \quad (15)$$

where N denotes the number of video samples.

For solution 1 introduced in Section 2.1, we begin by extracting the facial ROI regions using the landmark detection tool of OpenFace during the data pre-processing step. We then follow the setting described in [17], applying a sliding window size of 300 frames (10s) and a step size of 15 frames (0.5s) to generate MSTmap from the facial videos. For the spatial-temporal Transformer module, we set the dimensionality D to 128 and the number of layers L to 6. During pre-training, we use the AdamW optimizer with a learning rate of 1e-4 and a batch size of 4. Data augmentation techniques include random horizontal and vertical flipping as well as frequency up/down sampling are used. In the fine-tuning step with data labels, in addition to the self-supervised loss, we also add Negative Pearson Loss to further optimize the model. Besides, we use a smaller learning rate, *i.e.*, 1e-5, to finetune the model. For the VIPL-V2 dataset,

Table 1

The ablation study results of our solution 1 on the test dataset.

Pre-training	Fine-tuning	Loss	RMSE↓ (bpm)
UBFC-rPPG	VIPL-V2	$\mathcal{L}_{band} + \mathcal{L}_{sparse} + \mathcal{L}_{var}$	13.88440
		$\mathcal{L}_{band} + \mathcal{L}_{sparse} + \mathcal{L}_{var} + \mathcal{L}_{perio}$	12.30601
UBFC-rPPG + PURE	VIPL-V2	$\mathcal{L}_{band} + \mathcal{L}_{sparse} + \mathcal{L}_{var}$	11.52003
		$\mathcal{L}_{band} + \mathcal{L}_{sparse} + \mathcal{L}_{var} + \mathcal{L}_{perio}$	10.67180
UBFC-rPPG + PURE + MMSE-HR	VIPL-V2	$\mathcal{L}_{band} + \mathcal{L}_{sparse} + \mathcal{L}_{var}$	10.36720
		$\mathcal{L}_{band} + \mathcal{L}_{sparse} + \mathcal{L}_{var} + \mathcal{L}_{perio}$	9.93125

we split the training and validation subsets in a ratio of 8:2. For the HR estimation inference step, following previous work [3, 4], we apply a 1st-order Butterworth filter to convert the rPPG signal into an HR value with a cutoff frequency range of [0.66Hz, 3.0Hz], corresponding to [40, 180] beats per minute. Subsequently, we perform the PSD [27] to estimate HR for each video clip. For solution 2 elaborated in Section 2.2, we resample the videos to a frame rate of 30 and then perform face detection and cropping. We set the length of the video clip to 300 frames without overlapping. Following the setting in [11], the spatial resolution S is set to 2, and the sampled time interval Δt of each rPPG signal is set to 150 frames. Other settings are the same as solution 1.

For the ensemble strategy, we take the multiple best prediction results under different settings of both solution 1 and solution 2. Then we average the different predicted heart rates of each sample as the final result.

3.3. Experimental Results

Results for Solution 1. As shown in Table 1, we investigate the impact of different pre-training datasets and loss functions for solution 1. The results indicate that as the amount of pre-training data increases, the performance of the model improves accordingly. In our solution, we ultimately select the UBFC-rPPG [21], PURE [22], and MMSE-HR [23] datasets for pre-training. Additionally, we also investigate the impact of the proposed periodicity loss \mathcal{L}_{perio} . We can see that the incorporation of the periodicity loss consistently improves the performance of the model significantly across different settings. For instance, when the model is trained on the UBFC-rPPG, PURE, and MMSE-HR datasets, the introduction of the periodicity loss reduces RMSE from 10.35720 to 9.93125. This improvement underscores the effectiveness of the periodicity loss in mitigating abnormal periodic fluctuations in the predicted signal and maintaining temporal periodicity consistency.

Results for Solution 2. As shown in Table 2, we evaluate different pre-training datasets, loss functions, and model inputs to find the best setting for this task. Note that the DISFA dataset is a non-posed facial expression database. However, from the results, we can find that using it for pre-training can still achieve comparable performance. Apart from that, we can find the same conclusion as solution 1 that increasing the amount of pre-training datasets is beneficial to performance. In this solution, we choose DISFA, UBFC-rPPG, MMSE-HR, and PURE for pre-training. Additionally, we also evaluate different combinations of supervised

¹<https://www.kaggle.com/competitions/the-3rd-repps-t1/leaderboard>

Table 2

The ablation study results of our solution 2 on the test dataset. * denotes the normalized difference on model input.

Pre-training	Fine-tuning	Loss	RMSE↓ (bpm)
DISFA	VIPL-V2	$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt}$	11.81139
		$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt} + \alpha\mathcal{L}_{pear}$	12.01150
		$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt} + \beta\mathcal{L}_{mcc}$	11.29330
DISFA + MMSE-HR	VIPL-V2	$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt}$	11.35523
		$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt} + \alpha\mathcal{L}_{pear} + \beta\mathcal{L}_{mcc}$	10.72491
DISFA + UBFC-rPPG + MMSE-HR	VIPL-V2	$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt}$	10.37686
		$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt} + \beta\mathcal{L}_{mcc}$	11.03058
		$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt} + \alpha\mathcal{L}_{pear} + \beta\mathcal{L}_{mcc}$	10.75880
DISFA + UBFC-rPPG + MMSE-HR + PURE	VIPL-V2	$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt}$	10.62485
		$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt} + \beta\mathcal{L}_{mcc}$	10.19808
		$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt} + \alpha\mathcal{L}_{pear} + \beta\mathcal{L}_{mcc}$	11.01228
* DISFA + UBFC-rPPG + MMSE-HR + PURE	VIPL-V2	$\mathcal{L}_{pos}^{gt} + \mathcal{L}_{neg}^{gt} + \alpha\mathcal{L}_{pear} + \beta\mathcal{L}_{mcc}$	10.36316

Table 3

The results of the top-3 leaderboards on the test dataset in each challenge of RePSS. The best result is highlighted in **bold**, and the second-best result is underlined. The results of 1st and 2nd RePSS are provided by the report [28, 29], and the 3rd results are provided by the Kaggle competition page¹.

Team Name	Venue	Rank	Method Type	RMSE↓ (bpm)
Mixanik	1st RePSS	1	Supervised	10.68021
PoWeiHuang	1st RePSS	2	Supervised	14.16263
AWoyczyk	1st RePSS	3	Supervised	14.37509
Dr.L	2nd RePSS	1	Supervised	11.05
TIME	2nd RePSS	2	Supervised	11.44
The Anti-Spoofers	2nd RePSS	3	Supervised	14.51
Face AI	3rd RePSS	1	Self-supervised	8.50693
HFUT-VUT (Ours)	3rd RePSS	2	Self-supervised	<u>8.85277</u>
PCA_Vital	3rd RePSS	3	Self-supervised	8.96941

loss \mathcal{L}_s . The results show that both the time domain and frequency domain loss are helpful for model fine-tuning. Moreover, we evaluate the performance of normalized frame difference input, and it shows a comparable result with normal input. In the model ensemble phase, we added the frame difference-based manner as different feature forms.

Model Ensemble. In order to combine the advantages of Solution 1 and Solution 2, we use an ensemble strategy to integrate the best prediction results of these two solutions together. Specifically, we ensembled the models by taking the average value of the prediction results for Solution 1 and Solution 2, and then obtained the final prediction results. As shown in Table 3, we report the top-3 results on the test dataset for each RePSS challenge. Compared to other teams, we can see that our team achieves 2nd place, which is higher than the 3rd by 1.2%. This demonstrates that our proposed two self-supervision solutions can complementarily achieve more accurate and robust heart rate estimation. Compared to the results of the supervised methods in previous challenges, we can find that self-supervised methods improve performance

by a large margin. This indicates that self-supervised methods can capture rPPG-related signals from facial videos during the pre-train phase without requiring any real physiological signals.

4. Conclusion

In this paper, we present our solutions developed for self-supervised remote heart rate measurement of the 3rd RePSS challenge hosted at IJCAI 2024. Specifically, we propose two self-supervised HR estimation solutions that integrate spatial-temporal modeling and contrastive learning, respectively. By leveraging the ensemble strategy, our final submission takes second place with the RMSE score of 8.85277 bpm. In the future, we plan to address the issues in this challenge from other perspectives, *e.g.*, using video motion magnification algorithms [30] to capture the subtle change reflected in faces by heartbeats.

Acknowledgments

This work was supported by the National Key R&D Program of China (NO.2022YFB4500601), the National Natural Science Foundation of China (72188101,62272144,62020106007and U20A20183), the Major Project of Anhui Province(202203a05020011), and the Fundamental Research Funds for the Central Universities.

References

- [1] X. Li, J. Chen, G. Zhao, M. Pietikainen, Remote heart rate measurement from face videos under realistic situations, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2014, pp. 4264–4271.
- [2] X. Li, I. Alikhani, J. Shi, T. Seppanen, J. Junntila, K. Majamaa-Voltti, M. Tulppo, G. Zhao, 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), 2018, pp. 242–249.
- [3] W. Qian, D. Guo, K. Li, X. Zhang, X. Tian, X. Yang, M. Wang, Dual-path tokenlearner for remote photoplethysmography-based physiological measurement with facial videos, IEEE Transactions on Computational Social Systems (2024).
- [4] Q. Li, D. Guo, W. Qian, X. Tian, X. Sun, H. Zhao, M. Wang, Channel-wise interactive learning for remote heart rate estimation from facial video, IEEE Transactions on Circuits and Systems for Video Technology (2023).
- [5] X. Liu, B. Hill, Z. Jiang, S. Patel, D. McDuff, Efficientphys: Enabling simple, fast and accurate camera-based cardiac measurement, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2023, pp. 5008–5017.
- [6] S. Tang, R. Hong, D. Guo, M. Wang, Gloss semantic-enhanced network with online back-translation for sign language production, in: Proceedings of the 30th ACM International Conference on Multimedia, 2022, pp. 5630–5638.
- [7] J. Zhou, D. Guo, M. Wang, Contrastive positive sample propagation along the audio-visual event line, IEEE Transactions on Pattern Analysis and Machine Intelligence (2022).

- [8] K. Li, D. Guo, M. Wang, Vigt: proposal-free video grounding with a learnable token in the transformer, *Science China Information Sciences* 66 (2023) 202102.
- [9] D. Guo, K. Li, B. Hu, Y. Zhang, M. Wang, Benchmarking micro-action recognition: Dataset, methods, and applications, *IEEE Transactions on Circuits and Systems for Video Technology* (2024).
- [10] Y. Wei, Z. Zhang, Y. Wang, M. Xu, Y. Yang, S. Yan, M. Wang, Deraincyclegan: Rain attentive cyclegan for single image deraining and rainmaking, *IEEE Transactions on Image Processing* 30 (2021) 4788–4801.
- [11] Z. Sun, X. Li, Contrast-phys+: Unsupervised and weakly-supervised video-based remote physiological measurement via spatiotemporal contrast, *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2024) 1–18.
- [12] H. Lu, H. Han, S. K. Zhou, Dual-gan: Joint bvp and noise modeling for remote physiological measurement, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 12404–12413.
- [13] Z. Yu, W. Peng, X. Li, X. Hong, G. Zhao, Remote heart rate measurement from highly compressed facial videos: an end-to-end deep learning solution with video enhancement, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 151–160.
- [14] J. Speth, N. Vance, P. Flynn, A. Czajka, Non-contrastive unsupervised learning of physiological signals from video, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 14464–14474.
- [15] K. Li, J. Li, D. Guo, X. Yang, M. Wang, Transformer-based visual grounding with cross-modality interaction, *ACM Transactions on Multimedia Computing, Communications and Applications* 19 (2023) 1–19.
- [16] X. Niu, S. Shan, H. Han, X. Chen, Rhythmnet: End-to-end heart rate estimation from face via spatial-temporal representation, *IEEE Transactions on Image Processing* 29 (2019) 2409–2423.
- [17] X. Niu, Z. Yu, H. Han, X. Li, S. Shan, G. Zhao, Video-based remote physiological measurement via cross-verified feature disentangling, in: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II* 16, 2020, pp. 295–310.
- [18] J. Gideon, S. Stent, The way to my heart is through contrastive learning: Remote photoplethysmography from unlabelled video, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 3995–4004.
- [19] A. Bardes, J. Ponce, Y. Lecun, Vicreg: Variance-invariance-covariance regularization for self-supervised learning, in: *International Conference on Learning Representations*, 2022.
- [20] K. Zhang, Z. Zhang, Z. Li, Y. Qiao, Joint face detection and alignment using multitask cascaded convolutional networks, *IEEE Signal Processing Letters* 23 (2016) 1499–1503.
- [21] S. Bobbia, R. Macwan, Y. Benezeth, A. Mansouri, J. Dubois, Unsupervised skin tissue segmentation for remote photoplethysmography, *Pattern Recognition Letters* 124 (2019) 82–90.
- [22] R. Stricker, S. Müller, H.-M. Gross, Non-contact video-based pulse rate measurement on a mobile service robot, in: *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, 2014, pp. 1056–1062.

- [23] S. Tulyakov, X. Alameda-Pineda, E. Ricci, L. Yin, J. F. Cohn, N. Sebe, Self-adaptive matrix completion for heart rate estimation from face videos under realistic conditions, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2016, pp. 2396–2404.
- [24] S. M. Mavadati, M. H. Mahoor, K. Bartlett, P. Trinh, J. F. Cohn, Disfa: A spontaneous facial action intensity database, *IEEE Transactions on Affective Computing* 4 (2013) 151–160.
- [25] S. M. Mavadati, M. H. Mahoor, K. Bartlett, P. Trinh, Automatic detection of non-posed facial action units, in: 2012 19th IEEE International Conference on Image Processing, 2012, pp. 1817–1820.
- [26] X. Niu, H. Han, S. Shan, X. Chen, Vipl-hr: A multi-modal database for pulse estimation from less-constrained face video, in: Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part V 14, 2019, pp. 562–576.
- [27] P. Welch, The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms, *IEEE Transactions on Audio and Electroacoustics* 15 (1967) 70–73.
- [28] X. Li, H. Han, H. Lu, X. Niu, Z. Yu, A. Dantcheva, G. Zhao, S. Shan, The 1st challenge on remote physiological signal sensing (repss), in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 2020, pp. 314–315.
- [29] X. Li, H. Sun, Z. Sun, H. Han, A. Dantcheva, S. Shan, G. Zhao, The 2nd challenge on remote physiological signal sensing (repss), in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 2404–2413.
- [30] F. Wang, D. Guo, K. Li, M. Wang, Eulermormer: Robust eulerian motion magnification via dynamic filtering within transformer, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 38, 2024, pp. 5345–5353.