

RhythmNet: End-to-End Heart Rate Estimation From Face via Spatial-Temporal Representation

Xuesong Niu¹, Student Member, IEEE, Shiguang Shan¹, Senior Member, IEEE, Hu Han, Member, IEEE, and Xilin Chen, Fellow, IEEE

Abstract—Heart rate (HR) is an important physiological signal that reflects the physical and emotional status of a person. Traditional HR measurements usually rely on contact monitors, which may cause inconvenience and discomfort. Recently, some methods have been proposed for remote HR estimation from face videos; however, most of them focus on well-controlled scenarios, their generalization ability into less-constrained scenarios (e.g., with head movement, and bad illumination) are not known. At the same time, lacking large-scale HR databases has limited the use of deep models for remote HR estimation. In this paper, we propose an end-to-end RhythmNet for remote HR estimation from the face. In RhythmNet, we use a spatial-temporal representation encoding the HR signals from multiple ROI volumes as its input. Then the spatial-temporal representations are fed into a convolutional network for HR estimation. We also take into account the relationship of adjacent HR measurements from a video sequence via Gated Recurrent Unit (GRU) and achieves efficient HR measurement. In addition, we build a large-scale multi-modal HR database (named as VIPL-HR¹), which contains 2,378 visible light videos (VIS) and 752 near-infrared (NIR) videos of 107 subjects. Our VIPL-HR database contains various variations such as head movements, illumination variations, and acquisition device changes, replicating a less-constrained scenario for HR estimation. The proposed approach outperforms the state-of-the-art methods on both the public-domain and our VIPL-HR databases.

Index Terms—Remote heart rate estimation, rPPG, spatial-temporal representation, end-to-end learning.

Manuscript received March 29, 2019; revised August 26, 2019 and September 30, 2019; accepted September 30, 2019. Date of publication October 22, 2019; date of current version January 10, 2020. This work was supported in part by the National Key R&D Program of China under Grant 2017YFA0700800, in part by the Natural Science Foundation of China under Grant 61672496 and Grant 61702486, and in part by the External Cooperation Program of Chinese Academy of Sciences (CAS) under Grant GJHZ1843. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Jiwen Lu. (Corresponding author: Shiguang Shan.)

X. Niu and X. Chen are with the Key Laboratory of Intelligent Information Processing, Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China, and also with the School of Computer Science and Technology, University of Chinese Academy of Sciences, Beijing 100049, China (e-mail: xuesong.niu@vipl.ict.ac.cn; xlchen@ict.ac.cn).

S. Shan is with the Key Laboratory of Intelligent Information Processing, Institute of Computing Technology, Chinese Academy of Sciences (CAS), Beijing 100190, China, with the School of Computer Science and Technology, University of Chinese Academy of Sciences, Beijing 100049, China, and also with the CAS Center for Excellence in Brain Science and Intelligence Technology, Shanghai 200031, China (e-mail: sgshan@ict.ac.cn).

H. Han is with the Key Laboratory of Intelligent Information Processing, Institute of Computing Technology, Chinese Academy of Sciences (CAS), Beijing 100190, China, and also with the Peng Cheng Laboratory, Shenzhen 518055, China (e-mail: hanhu@ict.ac.cn).

Digital Object Identifier 10.1109/TIP.2019.2947204

¹VIPL-HR is available at: http://vipl.ict.ac.cn/view_database.php?id=15

1057-7149 © 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

I. INTRODUCTION

HEART rate (HR) is an important physiological signal that reflects the physical and emotional status of a person; therefore, HR measurement is useful in many applications, such as training aid, health monitoring, and nursing care. Traditional HR measurement usually relies on contact monitors, such as electrocardiograph (ECG) and contact photoplethysmography (cPPG) based sensors, which are inconvenient for the users and thus limit the application scenarios. In recent years, a growing number of studies have been reported on remote HR estimation from videos [1]–[7], which allows HR estimation from the skin, i.e., the face area, without contact with a person.

The existing video-based HR estimation methods can be grouped into two main categories: remote photoplethysmography (rPPG) based approaches [1], [2], [4]–[7] and the ballistocardiographic (BCG) [3], [11] based approaches. While rPPG-based approaches aim to extract HR signals from the color changes caused by variations in volume and oxygen saturation of the blood in the vessels due to heart beats, BCG-based approaches aim to extract HR signals from head movements arising from the periodic ejections of blood into the great vessels along with each heartbeat. Most of the existing non-contact HR measurement methods in literature are rPPG-based approaches because it is easier to capture the skin color changes than to get the small head movement in less-constrained scenarios, especially when the subjects perform large-scale head movements. However, head movements and environment conditions, e.g., light conditions, may have great impact on rPPG-based approaches since the amplitude of the HR signal extracted from the color changes of facial skin is too small to be perceived by human eyes.

In order to get stable HR signal, many rPPG-based methods have been proposed. Color space transformation and signal decomposition are commonly used tools for HR signal generation, e.g., using independent component analysis (ICA) [1], [2] and CHROM [4]. These methods reported good results under constrained situations, but may not perform well in the presence of large head motions and lighting variations. Some latter approaches introduced more sophisticated color space transformation and signal decomposition algorithms and obtained improved heart rate measurement accuracy [6], [7] under less-constrained situations. However, these methods usually made some certain assumptions w.r.t. the skin reflection (e.g., linear combination assumption [1], [2]), and the influence

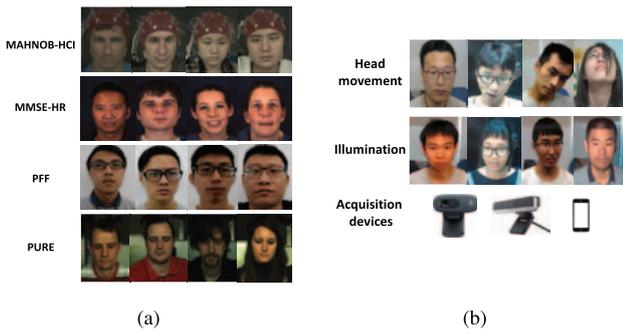


Fig. 1. Comparisons between (a) several public domain HR databases (MAHNOB-HCI [8], MMSE-HR [6], PFF [9], PURE [10]) and (b) our VIPL-HR database in terms of illumination condition, head movement, and acquisition device.

by head movement [4], [7]. Such assumptions may not hold in the presence of large head motions and lighting variations. In addition, many prior methods usually perform evaluations on private databases, making it difficult to provide comparisons among different methods. While a few public-domain HR databases are available [6], [8]–[10], [12], the sizes of these databases are quite limited (usually less than 50 subjects) compared with the large scale databases in the other computer vision tasks. Moreover, these databases were usually built in well-controlled scenarios, i.e., with minor illumination and motion variations (see Fig. 1(a)), which has a big gap with the practical application scenarios.

In this paper, we aim to perform remote HR estimation from face under less-constrained scenarios, where there are large variations in pose and illumination. In one aspect, we aim to build an end-to-end trainable HR estimator leveraging the strong modelling capacity of deep neural networks (DNN) to address the challenges of head movement, illumination variation, and acquisition device difference. In the other aspect, we also hope to provide an effective approach for modelling the relationship between adjacent HR measurement. To replicate practical scenarios, we also build a large-scale multi-modal HR database (named as VIPL-HR), which consist of 2,378 visible face videos and 752 NIR face videos from 107 subjects. Three different recording devices (RGB-D camera, smart phone, and web-camera) are used for video recording under 9 different situations covering variations in face pose, scale, and illumination (see Fig. 1). We provide evaluations on two widely used public-domain datasets (MAHNOB-HCI [8] and MMSE-HR [6]) and our VIPL-HR datasets covering both intra-database and cross-database testing protocols.

This work is an extension of our early work [13]. Compared with our previous work, the major extensions in this work are three-fold: i) a detailed review of the published methods for remote HR estimation, and in-depth analysis of the remaining challenges in remote HR estimation from face; ii) we take into account the correlation between adjacent HR measurements and model it via Gated Recurrent Unit (GRU) to achieve robust HR estimation; and iii) extensive evaluations are provided on multiple databases covering a number of aspects of intra-database testing, inter-database testing, visible light imaging,

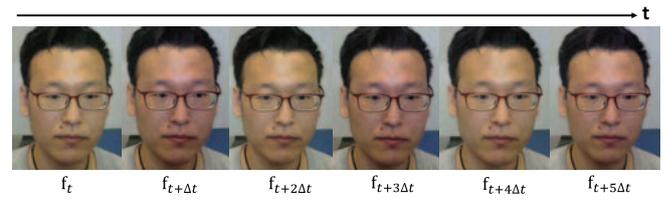


Fig. 2. Magnified skin color changes due to heartbeat using Eulerian Video Magnification [29].

near-infrared imaging, video compression, etc. The experimental results show that the proposed approach outperforms the state-of-the-art algorithms by a large margin.

The rest of this paper is organized as follows. We review the related work of remote HR estimation in Section II. The large-scale VIPL-HR database we built is introduced in Section III. The proposed end-to-end approach for remote HR estimation is detailed in Section IV. Experimental evaluations and analysis are provided in Section V. Finally, we conclude this work in Section VI.

II. RELATED WORK

A. Remote Heart Rate Estimation From Face Using rPPG

The principle of the rPPG-based HR estimation is based on the fact that optical absorption of skin varies periodically with the blood volume pulse (BVP). Human skin can be regarded as a three-layer model, subcutis, dermis and epidermis. The light is usually absorbed by the chromophores in the skin. The major chromophores of human skin are hemoglobin in the microvasculars of dermis and subcutis layers and melanin in the epidermis layer. The changes of hemoglobin content during a cardiac cycle would cause tiny color changes of the skin² (see Fig. 2). Although the color changes are invisible to human eyes, they can be captured by commodity RGB sensors and used for computing the HR [30], and a number of approaches have been reported to use such information for HR estimation.

Although it is possible to perform rPPG based HR measurement, it is still a challenge task since the amplitude of the light absorption variation w.r.t. BVP is minor. In addition, these signals can be easily affected by head movements, illumination conditions and noise introduced by the recording devices, making it hard to build a robust HR measurement system under less-constrained situations. A database considering the influences of these factors are also lacking to the community.

Blind signal separation (BSS) methods were introduced for remote HR estimation in [1], [2], in which independent component analysis (ICA) was applied to the temporal RGB color signal sequences to find the dominant component related to HR. In [14], ICA was used for color signals captured by a five-band camera instead of an RGB camera. Different from these approaches which applied ICA to the whole region of interest (ROI), Lam and Kuno applied ICA in a patch level for HR estimation [15]. Besides ICA, principal components analysis (PCA) [16] and constrained ICA (cICA) [17] were

²We provide a video (MagnificatedSkinColorChange.avi) of such magnified skin color changes in https://drive.google.com/file/d/1_sjPOaXP8FLQkGoJIPsDuU0mRSxegv/view?usp=sharing

also used to extract the purified HR signal from the input RGB video sequence. The BSS-based HR estimation approaches usually make certain assumptions about the HR signals. When these assumptions hold, BSS-based methods can effectively improve the signal-to-noise rate (SNR) of the HR signals. However, the assumptions may not hold under less-constrained scenarios, e.g., with large head movement, poor illumination, diverse sensors, etc. As a result, the HR estimation accuracy of BSS-based methods could degrade severely.

Since the rPPG-based methods is highly related to the skin, it is important to leverage the prior knowledge of skin to improve the robustness. An optical skin model considering head movement was proposed in [4], and used for calculating a chrominance feature representation to reduce the influence of head movement to HR estimation. In a later work of [18], pixel-wise chrominance features were computed and used for HR estimation. Instead of modelling the influence of head movement w.r.t. different color channels, Feng *et al.* [19] treated the human face as a Lambertian radiator and provided a radiance analysis of different motion types. Besides modelling the optical status of skin under head motion, Haan and Leest [20] took the camera optical properties into consideration and proposed a linear decomposition of RGB signals to perform HR estimation. A detailed discussion of different skin optical models for rPPG-based HR estimation was presented in [7], in which a new projection method was also proposed for extracting pulse signals from an input RGB video. The various skin models utilized by individual HR estimation methods lead to improved robustness. However, when it comes to some complicated scenarios, the assumptions used by these models may not hold, and the HR estimation accuracy could drop significantly. Models that can adaptively deal with various variations in practical application scenarios are required to build a robust HR estimation system. Thus, data-driven based feature learning methods may have the ability to obtain such a representation.

After obtaining the signals related to HR, a common operation is to apply Fast Fourier Transform (FFT) to the signals to get the spectral distribution. The peak of the spectral distribution is considered as the HR frequency (in Hz). Directly applying FFT to the HR signals may get a spectral distribution consisting of many spectral noises. Therefore, a few approaches have been proposed to reduce such noises. In [21], Kumar *et al.* used the frequency characteristics as the weights to combine the HR signals from different ROIs in order to get a better HR signal. Instead of enhancing the HR signals, Wang *et al.* [22] directly adjusted the spectral distribution according to the prior knowledge of the HR signal. Similarly, Niu *et al.* [23] took the continuous estimation situations into consideration and used preceding estimations to learn an HR distribution and use it to modulate the spectral distribution. While the frequency domain representation is helpful for remote HR estimation, temporal domain information can also be helpful.

All the methods mentioned above mainly focus on improving the hand-crafted pipelines from periodical temporal signal extraction to signal analysis. Their generalization ability can be poor under unconstrained scenarios. At the same time,

data-driven based methods are believed to have strong ability to learn more relevant features to a specific task, particularly when there are given enough data. Some recent methods tried to obtain better temporal signals of the HR from a face video. Tulyakov *et al.* [6] divided the face into multiple ROI regions to get a matrix of temporal representation and utilized matrix completion to purify the signals related to HR. Wang *et al.* [24] proposed a subject-dependent approach to compute the HR signal using spatial subspace rotation (SSR) given a complete continuous sequence of a subject's face. Another kind of data-driven methods for remote HR estimation is to learn a HR estimator using the frequency domain representations as input. Hsu *et al.* [25] combined the frequency domain features from RGB channels and the ICA components to get a HR signal representation and used support vector regression (SVR) to estimate the HR. Hsu *et al.* [9] generated time-frequency maps from the pre-processed green channel signals and used them as input of a VGG-16 model to estimate the HR. Chen and Mcduff [26] fed the original RGB face video into a convolutional network with attention mechanism and output the related HR signals. Although the existing data-driven approaches attempted to make use of statistical learning, as opposed to physical model based signal analysis, they failed to build an end-to-end HR estimator. In addition, the feature representations used by these methods remain hand-crafted, which may not be optimum for the HR estimation task.

Besides rPPG-based HR measurement, another kind of remote HR estimation methods is based on the ballistocardiographic (BCG) signals, which is the subtle head motions caused by cardiovascular circulation. Inspired by the Eulerian magnification method [29], Balakrishnan *et al.* performed facial key points tracking and used PCA to extract the pulse signal from the trajectories of key points [3]. Instead of using the PCA component of key points' trajectories to represent HR signal, Irani *et al.* [11] performed a stochastic search method using discrete cosine transform (DCT) to find the PCA component with the most significant periodicity. Since this kind of methods is based on subtle motion analysis, the voluntary movements by the subjects will significantly influence the HR estimation accuracy, leading to very limited use in practical applications.

Camera sensors can also have impact on the HR estimation accuracy. Besides the widely used commodity RGB sensors, several approaches also studied the possibility of using non-RGB, e.g., near-infra-red (NIR) sensors. Inspired by the PVB pulse-extraction method in [20], Gastel *et al.* investigated the feasibility of rPPG-based HR estimation in the NIR spectrum and achieved a robust HR estimation result with both NIR and color sensors [27]. Chen *et al.* [28] directly used the NIR images captured by RealSense³ and demonstrated the possibility of using NIR images for HR measurement. Instead of using non-color sensors to perform rPPG-based HR estimation, a BCG-based method using depth images was studied in [31]. Similar to the color sensor based methods, all the remote HR estimation using non-color sensors are still hand-crafted, and make certain assumptions.

³<https://www.intel.com/content/www/us/en/homepage.html>

TABLE I
A BRIEF SUMMARY OF THE EXISTING rPPG BASED REMOTE HR ESTIMATION METHODS

Publication	Camera	Input signal	HR temporal signal extraction	HR estimation	Results (RMSE in bpm)
Verkruijse et al. [14]	Webcam	G	bandpass filter	FFT	-
Poh et al. [1]	Webcam	R,G,B	ICA + bandpass filter	FFT	MAHNOB-HCI 25.9 bpm
Poh et al. [2]	Webcam	R,G,B	ICA + bandpass & detrending filter	Peak detection	MAHNOB-HCI 13.6 bpm
McDuff et al. [15]	Five-band camera	R,G,B,C,O	ICA + bandpass& detrending filter	Peak detection	-
Lam et al. [16]	Webcam	R,G,B	Patch level ICA + bandpass& detrending filter	FFT + majority vote	MAHNOB-HCI 8.9 bpm
Lewandowska et al. [17]	Webcam	R,G,B	PCA + bandpass filter	FFT	-
Tsouri et al. [18]	Webcam	R,G,B	constrained ICA + bandpass filter	FFT	-
Haan and Jeanne [4]	Webcam	R,G,B	Color space transformation based on skin model (CHROM feature)	FFT	MAHNOB-HCI / MMSE-HR / VIPL-HR 6.23 bpm / 11.37 bpm / 16.9 bpm
Wang et al. [19]	Webcam	R,G,B	Pixel level CHROM feature + PCA	FFT	-
Feng et al. [20]	Webcam	R,G	Adaptive Green/Red difference	FFT	-
Haan and Leest [21]	Webcam	R,G,B	Color space transformation based on skin-tone and camera properties	FFT	-
Wang et al. [7]	Webcam	R,G,B	Color space transformation based on skin model	FFT	VIPL-HR 17.2 bpm
Kumar et al. [22]	Webcam	G	Patch level signals combination using weighted average	FFT	-
Wang et al. [23]	Webcam	R,G,B	Using the prior knowledge of heart rate to remove noise	FFT	-
Niu et al. [24]	Webcam	R,G,B	Patch level CHROM feature	FFT + distribution learning	MAHNOB-HCI 8.72 bpm
Tulyakov et al. [6]	Webcam	R,G,B	CHROM features + self-adaptive matrix completion	FFT	MAHNOB-HCI / MMSE-HR / VIPL-HR 6.23 bpm / 11.37 bpm / 21.0 bpm
Wang et al. [25]	Webcam	R,G,B	Color space transformation based on spatial subspace of skin-pixels	FFT	-
Hsu et al. [26]	Webcam	R,G,B	ICA	FFT+SVR	-
Hsu et al. [9]	Webcam	R,G,B	CHROM feature + detrending filter + 2nd order difference	STFT + VGG model	MAHNOB-HCI 4.27 bpm
Chen et al. [27]	Webcam	R,G,B	Convolutional Network with attention module	FFT	MAHNOB-HCI (MAE) 4.57 bpm
Li et al. [5]	Webcam	G	Background noise removing + bandpass & detrending filter	FFT	MAHNOB-HCI / MMSE-HR 7.62 bpm / 19.95 bpm
Van et al. [28]	NIR Sensor	NIR	Signal modified using prior of skin reflectance	FFT	-
Chen et al. [29]	RealSense	NIR	Global self-similarity filter	FFT	-
Proposed	Webcam/Phone RealSense	R,G,B/NIR	Spatial-temporal maps	Deep regression model with GRU	MAHNOB-HCI / MMSE-HR / VIPL-HR 4.00 bpm / 5.03 bpm/ 8.14 bpm

The existing remote HR measurement methods are summarized in Table I. Despite tremendous progress has been made in remote HR estimation, there are still limitations. First,

the existing approaches usually made certain assumptions when building the model, which may limit the application scenarios. Second, most of the existing approaches are designed

TABLE II
A SUMMARY OF THE PUBLIC-DOMAIN DATABASES AND OUR VIPL-HR DATABASE FOR REMOTE HR ESTIMATION

	# Subjects	Illumination variation	Head movement	Sensor diversity	#Videos
MAHNOB-HCI [8]	27	<i>L</i>	<i>E</i>	<i>C</i>	527
MMSE-HR [6]	40	<i>L</i>	<i>E</i>	<i>C</i>	102
PURE [10]	10	<i>L</i>	<i>S/SM/T</i>	<i>C</i>	60
PFF [9]	13	<i>L/D</i>	<i>S/SM</i>	<i>C</i>	104
OBF [12]	106	<i>L</i>	<i>S</i>	<i>C/N</i>	2,120
VIPL-HR	107	L/D/B	S/LM/T	C/N/P	3,130

L = Lab Environment, D = Dim Environment, B = Bright Environment, E = Expression,
S = Stable, SM = Slight Movement, LM = Large Movement, T = Talking,
C = Color Camera, N = NIR Camera, P = Smart Phone Frontal Camera

in a step-by-step hand-crafted way requiring sophisticated knowledge and experiences. Therefore, end-to-end trainable approaches are required to bridge the gap between the capability of remote HR estimation methods and the real application requirements.

B. Public Databases for Remote HR Estimation

Many existing methods reported their performance using private databases, leading to difficulties in performance comparison across individual approaches. Public-domain database for evaluating remote HR measurement methods are quite limited. MAHNOB-HCI database [8] was firstly used for remote HR estimation in [5], which consists of 527 videos of 27 subjects, recorded with small head movement and facial expression variation under laboratory illumination. MMSE-HR [6] was introduced by Tulyakov *et al.*, which consists of 102 videos of 40 subjects involving various subjects' facial expressions and head movements. These two databases are originally designed for emotion analysis, and the subjects' variousness is mainly limited to facial expression changes. At the same time, all the videos of MAHNOB-HCI and MMSE-HR are compressed, which may cause some damages to the HR signals used for further calculation.

There are also a few public-available databases especially designed for the task of remote HR estimation. Stricker *et al.* [10] released the PURE database consisting of 60 videos from 10 subjects, in which all the subjects are asked to perform 6 kinds of movements such as talking or head rotation. Hus *et al.* [9] released the PFF database, consisting of 104 videos of 10 subjects, in which only small head movements and illumination variations are involved. These two databases are limited by the number of subjects and recording scenarios, making them not suitable for building a real-world HR estimator. In 2018, Xiaobai *et al.* [12] proposed the OBF database specifically designed for heart rate variability (HRV) analysis, and all the scenarios in this database are well-controlled, making it very easy to measure HR.

We summarize the public-domain databases for remote HR estimation in Table II. We can see that all the existing databases are limited in either the number of subjects or the recording scenarios, particularly limited to constrained scenarios. A large-scale database recorded under less-constrained environment is required for studying remote HR estimation methods in practice.

III. VIPL-HR DATABASE

Considering the constrained data acquisition scenarios used by most of the existing databases for remote HR estimation, we built a new dataset, named as VIPL-HR to promote the research of remote HR estimation under less-constrained scenarios, *i.e.*, with head movement, illumination variation, and acquisition device diversity. VIPL-HR contains 3,130 face videos from 107 subjects, which are recorded with three different sensors (webcam, RealSense, and smartphone) under varying illumination and pose variations. In this section, we provide the details in building our VIPL-HR database covering: i) device setup and data collection, ii) face video compression, and iii) database statistics.

A. Device Setup and Data Collection

We design our data collection procedure with the objective in mind that face video recording conditions should cover diversities of environmental illumination, subject pose, acquisition sensor, and HR distribution in order to replicate daily application scenarios.

The recording environmental setup is illustrated in Figure 3. In order to cover diverse illumination conditions, we utilize a filament lamp placed in front of the subject as well as the ceiling lamp. Three different illumination scenarios are conducted, *i.e.*, both lamps are turned on, only the ceiling lamp is turned on, and both lamps are turned off. In order to cover diverse pose variation of the subject, we first ask the subjects to sit casually in front of the cameras, then encourage them to perform daily activities such as talking and looking around. In order to cover a wider range of HR in our VIPL-HR database, we ask each subject to do some exercises and then came back for database recording. At the same time, different distances between the subject and the cameras are also considered.

Sensor diversity is also considered when building the database. We choose three different cameras, *i.e.*, Logitech C310 webcam, Intel Realsense F200 RGB-D camera, and Huawei P9 smartphone (with its frontal camera). Logitech C310 and RealSense F200 are used for visible face video recording with different video resolution, and RealSense F200 is also used for NIR face videos recording to investigate the possibility of remote HR estimation under dark lighting conditions. Since smartphones have become

TABLE III
SPECIFICATIONS OF INDIVIDUAL RECORDING DEVICES USED IN OUR VIPL-HR DATABASE

Device	Specification	Setting	Output
Computer	Lenovo ThinkCentre	Windows 10 OS	N/A
Color camera	Logitech C310	25fps*, 960×720 color camera	Color videos
RGB-D camera	RealSense F200	30fps*, 640×480 NIR camera 1920×1080 color camera,	Color videos NIR videos
Smart phone	HUAWEI P9 frontal camera	30fps, 1920×1080 color camera	Color videos
BVP recoder	CONTEC CMS60C	N/A	HR, SpO2, and BVP signals
Filament lamp	N/A	50Hz	N/A

* The frame rates of Logitech C310 and RealSense F200 are not very stable, and the time steps corresponding to each frame are recorded.

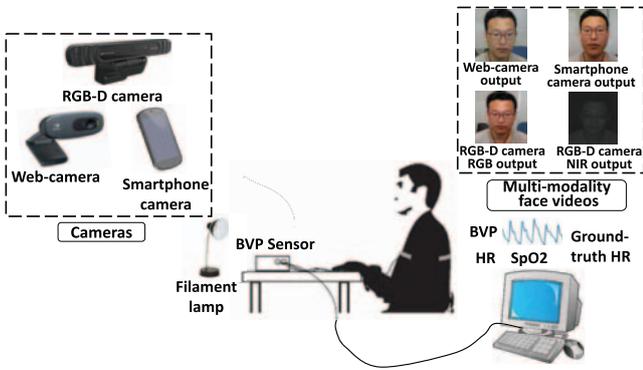


Fig. 3. Devices and setup used in collecting VIPL-HR.

an indispensable part of our daily lives, we also record the face videos when the smartphone is placed in front of the subjects. At the same time, the subjects are also asked to hold the smartphone by themselves to record videos like a video chat scenario. All the related physiological signals, including HR, SpO2, and BVP signals, are synchronously recorded with a CONTEC CMS60C BVP sensor. The details of these recording scenarios can be found in Table IV and the details of the device specifications can be found in Table III.

B. Face Video Compression

The raw data of our VIPL-HR is about 1.05 TB in total, making it inconvenient for distribution. In order to facilitate access to our dataset by the other researchers, we investigate to make a compressed and resized version of our database. As studied in [32], video compression may have big influence on video-based heart rate estimation. Therefore, we carefully compare the individual video compress methods, and frame resizing methods.

We considered five different video codecs, i.e., ‘MJPG’, ‘FMP4’, ‘DIVX’, ‘PIM1’ and ‘X264’,⁴ which are commonly used in the video coding community. The resizing scales we considered are 1/2, 2/3, and 3/4 for each dimension of the

⁴<http://www.fourcc.org/>

TABLE IV
DETAILS OF THE NINE RECORDING SITUATIONS
IN THE VIPL-HR DATABASE

Scenario	Head movement	Illumination	Distance	Exercise	Phone recording method
1	S	L	1m	No	Fixed
2	LM	L	1m	No	Fixed
3	T	L	1m	No	Fixed
4	S	B	1m	No	Fixed
5	S	D	1m	No	Fixed
6	S	L	1.5m	No	Fixed
7	S	L	1m	Yes	Fixed
8	S	L	-	No	Hand-held
9	LM	L	-	No	Hand-held

S = Stable, LM = Large Movement, T = Talking,
L = Lab Environment, D = Dim Environment, B = Bright Environment

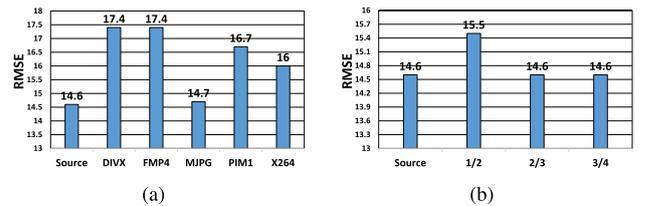


Fig. 4. Evaluations on VIPL-HR using (a) different video compression codecs and (b) different image resolutions for data compression. All tests are based on Haan2013 [4].

original frame. All the videos are compressed using OpenCV⁵ with FFMpeg interface.⁶ We choose a widely used remote HR estimation method Haan2013 [4] as a baseline HR estimation method to verify how data compression influence its HR estimation accuracy.

The HR estimation accuracies in terms of root mean square error (RMSE) by the baseline HR estimator [4] on various compressed videos are reported in Fig. 4. From the results, we can see that the ‘MJPG’ video codec works better in maintaining the HR signal in the videos while it is able to reduce

⁵<https://opencv.org/>

⁶<http://ffmpeg.org/>

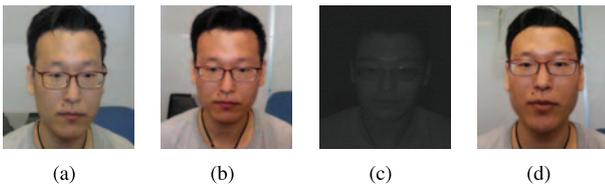


Fig. 5. Example video frames captured by different devices: (a) Logitech C310, (b) RealSense F200 color camera, (c) RealSense F200 NIR camera, and (d) HUAWEI P9 frontal camera.

the size of the database significantly. Resizing the frames to less than two-thirds of the original image resolution leads to damage to the HR signal. Therefore, we choose the ‘MJPG’ codec and 2/3 resizing as our final data compression solution. Finally, we obtained a compressed VIPL-HR dataset with about 48 GB. We would like to share both the uncompressed and compressed databases to the research community based on individual researchers’ preference.

C. Database Statistics

The VIPL-HR dataset contains a total of 2,378 color videos and 752 NIR videos from 107 participants with 79 males and 28 females, and ages distributed between 22 and 41. Each video is recorded with a length of about 30s, and the frame rate is about 30fps (see Table II). Some example video frames of one subject captured by different devices are shown in Fig. 5.

To obtain additional statics about the pose variations in our VIPL-HR database, we estimate head poses (yaw, pitch, and roll angles) for each video using the OpenFace head pose estimator⁷ (see Situation 2 in Table IV). Histograms for maximum amplitudes of yaw, pitch, and roll rotations for all the videos can be found in Fig. 6. From the histograms, we can see that the maximum rotation amplitudes of the subjects vary in a large range, i.e., the maximum rotation amplitudes are 92° in roll, 105° in pitch and 104° in yaw. This is reasonable because every subject is allowed to look around during the video recording process.

At the same time, in order to quantitatively demonstrate the illumination changes in VIPL-HR database, we have calculated the mean gray-scale intensity of face area for Situation 1, Situation 4, and Situation 5 in Table IV. The results are shown in Fig. 7. We can see that the mean gray-scale intensity varies from 60 to 212, covering complicated illumination variations.

A histogram of ground-truth HRs is also shown in Fig. 8. We can see that the ground-truth HRs in VIPL-HR vary from 47 bpm to 146bpm, which covers the typical HR range.⁸ The wide HR distribution in VIPL-HR mitigates the gap between the HR range covered in dataset collection and the HR distribution that can appear in daily-life scenes. The relatively large size of our VIPL-HR also makes it possible to leverage deep learning methods to build more robust data-driven HR estimation model.

IV. PROPOSED APPROACH

Prior work performed remote HR estimation mainly based on hand-crafted features, which are built based on certain

⁷<https://github.com/TadasBaltrusaitis/OpenFace>

⁸https://en.wikipedia.org/wiki/Heart_rate

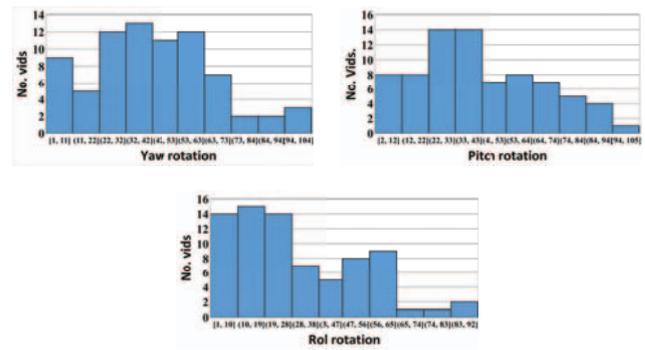


Fig. 6. Histograms of the maximum amplitudes (in degree) of the yaw, pitch, and roll rotations for all the videos with head movement in VIPL-HR.

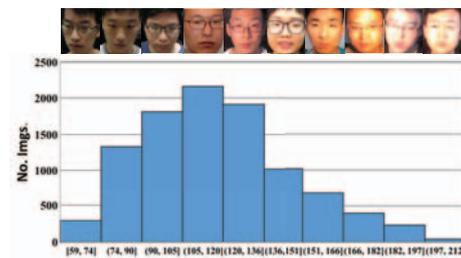


Fig. 7. The histogram of the average image intensity (gray-scale) of the videos recorded under the illumination-specific situations in VIPL-HR.

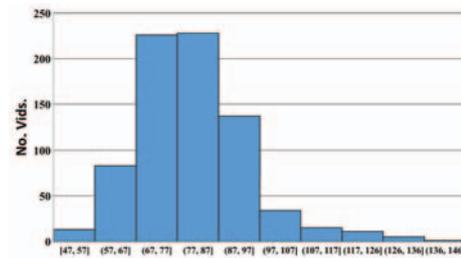


Fig. 8. The histogram of the ground-truth HR distribution in VIPL-HR.

assumptions of the HR signal and environment. These methods may fail in realistic situations when the assumptions do not hold. Given the large-scale VIPL-HR database with diversities in illumination, pose, and sensors, we are able to build a data-driven HR estimation model using deep learning methods. In this section, we are going to introduce the RhythmNet for remote HR estimation from face video. Fig. 9 gives an overview of our RhythmNet. We first perform detection and facial landmark detection to locate region of interest (ROI) facial areas, from which a spatial-temporal map is computed as a novel low-level representation of the HR signal. Finally, a CNN-RNN model is designed to learn a mapping from the spatial-temporal maps to the HR.

A. Face Detection, Landmark Detection and Segmentation

We use an open source face detector SeetaFace⁹ to detect the face and localize 81 facial landmarks (see Fig. 10). Since face

⁹<https://github.com/seetaface/SeetaFaceEngine>

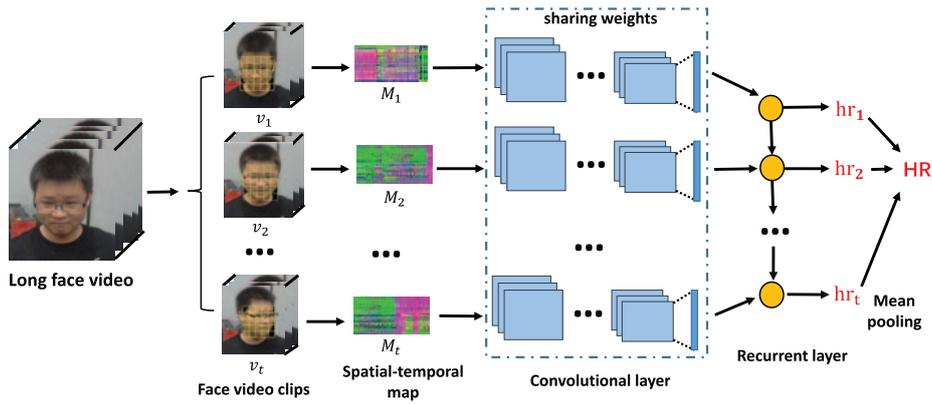


Fig. 9. Overview of our RhythmNet. Given an input video sequence, we first divide it into multiple short video clips (v_1, v_2, \dots, v_t), and perform face and landmark detection on each frame for face alignment. Then spatial-temporal maps are generated from the aligned face images per video clip to represent the HR signals, and a deep network consisting of convolutional layers and recurrent layers is trained to predict the HR from the spatial-temporal maps. Finally, the HR estimated for the input video sequence is computed as the average of all the estimated HRs from individual video clips.

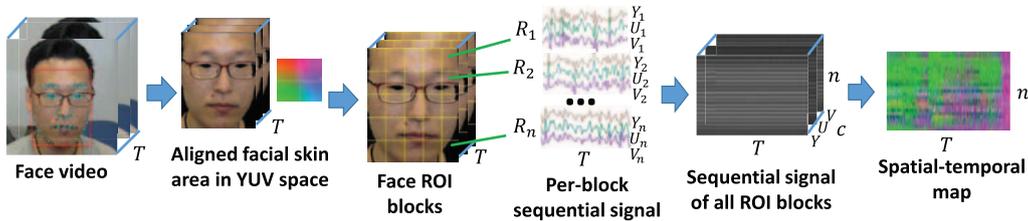


Fig. 10. An illustration of spatial-temporal map generation from face video. We first align faces across different frames based on the detected facial landmarks, and transform the aligned face images to YUV color space. Then the facial area is divided into n ROI blocks R_1, R_2, \dots, R_n , and the average color value is computed for each color channel within each are concatenated into a sequences, i.e., $Y_1, U_1, V_1, Y_2, U_2, V_2, \dots, Y_n, U_n, V_n$. These $n * c$ sequences are placed into rows to form a spatial-temporal map with the size of $T \times n \times c$.

and facial landmarks detection is able to run at a frame rate of more than 30 fps, we perform face detection and landmarks detection on every frame in order to get accurate and consistent ROI facial areas in a video sequence. A moving average filter is also applied to the 81 facial landmarks across frames to get more stable landmark locations.

In order to make full use of all the informative parts containing the color changes due to heart rhythms, we choose to use the whole face area for further processing. Face alignment is firstly performed using the eye centre points, and then a face bounding box is defined with a width of w (where w is the horizontal distance between the outer cheek border points) and height $1.2 * h$ (where h is the vertical distance between chin location and eye eyebrow centre). Skin segmentation is then applied to the defined ROI to remove the non-face area such as eye region and background area. Since resizing may introduce noise to the HR signals, we choose to use the original face region cropped based on facial landmarks for further processing.

B. Spatial-Temporal Map for Representing HR Signals

The **only** useful information for rPPG-based HR estimation in face video is the skin color variations caused by variations in volume and oxygen saturation of the blood in the vessels due to heart beats, which are very minor in amplitude and can be affected by head movements, illumination variations

and sensor noises. In order to suppress noises and improve the SNR of HR signals, Most rPPG-based HR estimation methods use the average pixel values of RGB channels of the whole face as the HR signal representation [1], [2], [4], [7], [19]. Such an average pooling representation provides better robustness than each single pixel. This can be regarded as the empirical rule, and can work smoothly for rPPG-based HR estimation under various conditions. Unlike the existing methods which only use the average pooling of ROI blocks to compute the HR signal, we have proposed a spatial-temporal map (see Fig. 10) to highlight the heart rhythm signals while suppressing the other information irrelevant to heart rhythm. Such a spatial-temporal map can enforce the input to the succeeding network to be as specific as to the heart rhythm signal, and make it possible for us to leverage CNN to learn informative representation for the final HR estimation task.

Specifically, for a video clip with T frames and c color space dimensions, we first get the face areas for all the frames as stated in IV-A. Then, the face area in each frame is divided into n ROI blocks R_1, R_2, \dots, R_n . The amplitude of these skin color variations are very weak, and average pooling has been proven to be effective in removing noises in rPPG-based HR estimation methods. Let $C(x, y, t)$ denote the value at location (x, y) of the t^{th} frame from different dimensions of the color space, and the average pooling of the i^{th} ROI block

of the t^{th} frame for each channel can be stated as

$$\bar{C}_i(t) = \frac{\sum_{x,y \in ROI_i} C(x,y,t)}{|ROI_i|} \quad (1)$$

where $|ROI_i|$ denotes the area of a block ROI (the number of pixels). So, for each video clip we can obtain $3 \times n$ temporal sequences with the length of T for different dimensions of the color space, e.g., $\mathbf{C}_i = \{\bar{C}_i(1), \bar{C}_i(2), \dots, \bar{C}_i(T)\}$, where C donates one of the c color space dimensions and i donates the index of the ROI. In order to make the best use of the HR signals, a min-max normalization is applied to each temporal signal, and the values of the temporal series are scaled into $[0, 255]$. Then, we place the n temporal sequences into rows and obtain a spatial-temporal map representation for the HR signal in a video clip. Eventually, we get a spatial-temporal representation from the raw face video clip with the size of $T \times n \times c$ as the input of our succeeding deep HR estimation network.

Color space selection is very important for representing the HR signal. As stated in [33], alternative color spaces derived from RGB video are beneficial for getting a better HR signal representation. Therefore, we take two kinds of color spaces into consideration: i) color spaces with dimension directly related to color, such as HSV color space, and ii) color spaces which are commonly used for skin segmentation, i.e., YUV and YCrCb color spaces. Testing on different color spaces can be found in Section V-D. We imperially choose to use the YUV color space (see our evaluations in Section V-D). The color space transformation can be formulized as

$$\begin{bmatrix} Y \\ U \\ V \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ -0.169 & -0.331 & 0.5 \\ 0.5 & -0.419 & -0.081 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} + \begin{bmatrix} 0 \\ 128 \\ 128 \end{bmatrix} \quad (2)$$

Another situation we need to consider is face detection failures for a few frames, which may happen when the subject's head moves or rotates too fast. This will lead to missing data of HR signals and the spatial-temporal representation. In order to handle this issue, we randomly mask a small part of the spatial-temporal maps along the time dimension to simulate the missing data cases, and use the partially masked spatial-temporal maps as augmented data to enhance the robustness of our RhythmNet.

C. Temporal Modeling for HR measurement

For each face video, we divide it into individual video clips, i.e., v_1, v_2, \dots, v_t , using a fixed sliding window containing w frames moved with a step of 0.5 seconds. The video clips are then used for computing the spatial-temporal maps of each clip, which are used as the input to our succeeding network. We choose ResNet-18 [34] as the backbone convolutional layers, which includes four blocks made up of convolutional layers and residual link, one convolutional layer, and one fully connected layer for the final regression. L_1 loss is used for measuring the difference between the predicted HR and ground truth HR. The output of the network is a single HR value regressed by a fully connected layer. All the estimated HRs

from individual video clips are normalized based on the frame rate of the face video.

We further take the relationship between adjacent measurements of two video clips into consideration and utilize a Gated Recurrent Unit (GRU) [35], [36] consisting of a cell, a reset gate and an update gate to model the temporal relationship between succeeding measurement. To be specific, the features extracted from the backbone CNN are fed to a one-layer GRU structure. The output of GRU is fed into a fully connected layer to regress the HR values for individual video clips. For each face video, the average of all the predicted HRs for individual video clips are computed as the final HR result.

Since the variance of the subjects' HRs is small during a very small period of time [23], we introduce a smooth loss function to constrain the smoothness of adjacent HR measurements. We take T continuous measurements hr_1, hr_2, \dots, hr_T into consideration, which are estimated from the GRU for a small period of time, i.e., 3 seconds. The mean HR within the T continuous measurements is then computed as $hr_{mean} = \frac{1}{T} \sum_{i=1}^T hr_i$ and the smooth loss L_{smooth} is defined as

$$L_{smooth} = \frac{1}{T} \sum_{i=1}^T \|hr_i - hr_{mean}\| \quad (3)$$

During the backpropagation, the partial derivative of the smooth loss L_{smooth} with respect to the input hr_t can be computed as

$$\frac{\partial L_{smooth}}{\partial hr_t} = \left(\frac{1}{T} - 1\right) sgn(hr_{mean} - hr_t) + \sum_{i=1, i \neq t}^T \frac{1}{T} sgn(hr_{mean} - hr_i) \quad (4)$$

where $sgn(x)$ is a sign function. The final loss function can be written as

$$L = L_{l1} + \lambda L_{smooth} \quad (5)$$

where L_{l1} denotes the L_1 loss function and λ is a parameter for balancing the two terms.

While RNN is widely used for modeling the temporal continuity in object tracking, action recognition, etc., its effectiveness in modeling the temporal relationships of physiological signals (such as HR), is not known. Therefore, our work is the first known approach uses RNN to improve the temporal estimation stability of HR.

V. EXPERIMENTS

In this section, we provide evaluations of the proposed RhythmNet covering three aspects: i) intra-database testing, ii) cross-database testing, and iii) key components analysis.

A. Database and Experimental Setting

We provide evaluations on two widely used public-domain databases, i.e., MAHNOB-HCI [8] and MMSE-HR [6] as well as the VIPL-HR database we collected. The details about these databases can be found in Table II. The ground-truth HRs of MAHNOB-HCI and MMSE-HR database are computed

from the electrocardiography (ECG) signals provided in the databases using the OSET ECG Toolbox.¹⁰

Different metrics have been used in the literature for evaluating the HR estimation performance, such as the mean and standard deviation (Mean and Std) of the HR error, the mean absolute HR error (MAE), the root mean squared HR error (RMSE), the mean of error rate percentage (MER), and Pearson's correlation coefficients r [5], [6]. We also use these evaluation metrics in our experiments below.

For the MMSE-HR and VIPL-HR datasets, we use a temporal sliding window with 300 frames to compute the spatial-temporal maps. For the MAHNOB-HCI database, since its original frame rate is 61 fps, we first downsample the videos to 30.5 fps before using the same temporal window size. The number of the ROI blocks used for spatial-temporal map generation for all the database is 25 (5×5 grids). During the training phase, half of the generated spatial-temporal maps were randomly masked, and the mask length varies from 10 frames to 30 frames. For the temporal relationship modelling, six adjacent estimated HRs are used to compute the L_{smooth} . The balance parameter λ in Equ. 5 is set to 100. Since the face sizes in the NIR videos are small for the face detector, only 497 NIR videos with detected faces are involved in the experiments. Our RhythmNet is implemented using PyTorch.¹¹ We use an Adam solver [37] with an initial learning rate of 0.001, and set the maximum epoch number to 50.

B. Intra-Database Testing

We first conduct intra-database experiments on each of the three databases. For the intra-database testing, in order to prove the effectiveness of the proposed HR estimator, we use a five-fold **subject-exclusive** cross-validation for the VIPL-HR database, which means that the subjects in the training set will not appear in the testing set. Since the MAHNOB-HCI and MMSE-HR databases have very limited face videos, we use a three-fold subject-independent cross-validation protocol.

1) *Experiments on Color Face Videos:* We first perform intra-database evaluations using RGB face videos in our large VIPL-HR database, and compare the proposed approach with a number of state-of-the-art traditional methods (Haan2013 [4], Tulyakov2016 [6], POS [7]). Two deep learning based methods, i.e., I3D [38] and DeepPhy [26], are also used for comparison. I3D [37] is widely used for video analysis, and DeepPhy [27] is the state-of-the-art deep learning based HR estimation method. For I3D, we directly use the code provided by DeepMind,¹² and trained it for HR estimation task. For DeepPhy, we contacted the authors, but the code is not publicly available; so we have re-implemented DeepPhy based on the descriptions in [27]. We use the same five-fold subject-exclusive protocol for all the deep learning based methods. For our RhythmNet, we report both results using and without using GRU for comparisons with the traditional state-of-the-art methods. The results of individual HR estimation methods are given in Table V.

¹⁰<http://www.oset.ir>

¹¹<https://pytorch.org/>

¹²<https://github.com/deepmind/kinetics-i3d>

TABLE V
THE HR ESTIMATION RESULTS BY THE PROPOSED APPROACH AND SEVERAL STATE-OF-THE-ART METHODS ON COLOR FACE VIDEOS OF THE VIPL-HR DATABASE

Method	Mean (bpm)	Std (bpm)	MAE (bpm)	RMSE (bpm)	MER	r
Tulyakov2016 [6]	10.8	18.0	15.9	21.0	26.7%	0.11
POS [7]	7.87	15.3	11.5	17.2	18.5%	0.30
Haan2013 [4]	7.63	15.1	11.4	16.9	17.8%	0.28
I3D [38]	1.37	15.9	12.0	15.9	15.6%	0.07
DeepPhy [27]	-2.60	13.6	11.0	13.8	13.6%	0.11
RhythmNet w/o GRU	1.02	8.88	5.79	8.94	7.38%	0.73
RhythmNet	0.73	8.11	5.30	8.14	6.71%	0.76

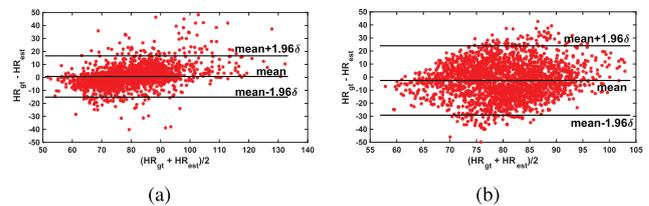


Fig. 11. The Bland-Altman plots demonstrating the agreement of the HR_{est} and HR_{gt} for (a) RhythmNet and (b) DeepPhy [26] in the VIPL-HR database. The lines represent the mean and 95% limits of agreement.

From Table V, we can see that the proposed approach achieves a promising result with a RMSE of 8.94 bpm when only using the CNN part of RhythmNet, which is a much lower error than the best of the baseline methods [26] (RMSE = 13.8 bpm). The HR estimation error by the complete RhythmNet is further reduced to 8.14 bpm. In Fig. 11(a), we also draw a Bland-Altman plot of the ground-truth HR and the estimated HR by our RhythmNet to analyze the estimation consistencies by individual approaches. The Bland-Altman plot for DeepPhy [26] is also given for comparison in Fig. 11(b). Again, it can be seen that our method achieves a better consistency than DeepPhy [26] on the VIPL-HR database.

We further check the HR estimation error distributions of the proposed method and [26]. As shown in Fig. 12, for most of the samples (71%), the proposed approach achieves a lower HR estimation error than 5 bpm, while the percentage for [26] is only 41.5%. We also plot the estimated HR comparing against the ground-truth HR in Fig. 13 to see the correlations between the estimated HR and the ground-truth HR. We can see that overall the predicted HRs are well correlated with the ground truth in a wide range of HR from 47 bpm to 147 bpm. The results show the effectiveness of the proposed end-to-end learning-based HR estimation method in handling challenges due to pose, illumination, and sensor variations.

We can also notice from Table V that deep learning based methods (our RhythmNet, I3D [38], and DeepPhy [26]) outperform the traditional methods (e.g., Tulyakov2016 [6], POS [7], Haan2013 [4]) on the challenging VIPL-HR dataset. These results indicate that given a large-scale database (such as VIPL-HR), deep learning based methods are able to learn more informative representation for HR estimation.

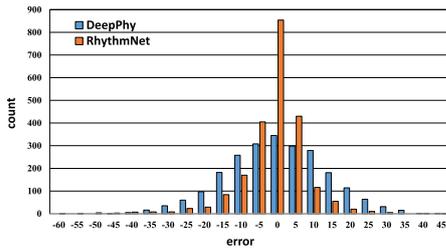


Fig. 12. Comparison of the HR estimation error distributions between the proposed RhythmNet and a learning-based state-of-the-art method DeepPhy [26].

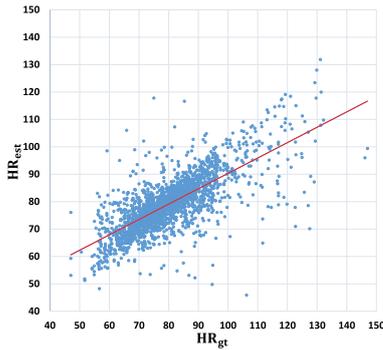


Fig. 13. The scatter plot comparing the ground truth HR_{gt} and the estimated HR_{est} by our RhythmNet for RGB face videos on the VIPL-HR dataset.

TABLE VI
THE HR ESTIMATION RESULTS ON NIR FACE VIDEOS
OF OUR VIPL-HR DATABASE

Method	Mean (bpm)	Std (bpm)	MAE (bpm)	RMSE (bpm)	MER	r
RhythmNet w/o GRU	1.26	13.8	9.11	13.9	11.6%	0.62
RhythmNet	1.74	12.4	8.45	12.5	10.8%	0.71

We also perform intra-database testing on the MAHNOB-HCI and MMSE-HR databases using a three-fold subject-exclusive protocol as stated in Section V-A. The results can be found in Table VII and Table VIII, respectively. From the results, we can see that the proposed RhythmNet again achieves the best results on both databases. The results indicate that our method has a good generalization ability under different image acquisition conditions, and is fairly robust when the training dataset become small.

2) *Experiment on NIR Face Video*: The experiments on NIR face videos are also conducted using the five-fold subject-exclusive protocol. Since the NIR face video only has one channel, we do not use color space transformation and compute one-channel spatial-temporal map for our RhythmNet. The baseline methods [4], [6], [7], [26] for RGB based HR estimation do not work on NIR videos. Thus, we only report the results by our RhythmNet without GRU and the complete RhythmNet in Table. VI. The Bland-Altman plots for NIR face videos is also given in Fig. 14.

From Fig. 11(a) and Fig. 14, we can see that the results using NIR face videos are not as good as those using RGB

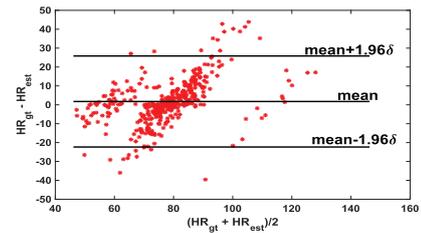


Fig. 14. The Bland-Altman plots demonstrating the agreement of the estimated HR_{est} by RhythmNet and the ground-truth HR_{gt} for NIR face videos in VIPL-HR. The lines represent the mean and 95% limits of agreement.

TABLE VII
THE HR ESTIMATION RESULTS BY THE PROPOSED APPROACH AND
SEVERAL STATE-OF-THE-ART METHODS ON THE
MAHNOB-HCI DATABASE

Method	Mean (bpm)	Std (bpm)	RMSE (bpm)	MER	r
Poh2010 [1]	-8.95	24.3	25.9	25.0%	0.08
Poh2011 [2]	2.04	13.5	13.6	13.2%	0.36
Balakrishnan2013 [3]	-14.4	15.2	21.0	20.7%	0.11
Li2014 [5]	-3.30	6.88	7.62	6.87%	0.81
Haan2013 [4]	-2.89	13.67	10.7	12.9%	0.82
Tulyakov2016 [6]	3.19	5.81	6.23	5.93%	0.83
RhythmNet(WithinDB)	0.41	3.98	4.00	4.18%	0.87
RhythmNet(CrossDB)	-5.66	6.06	8.28	8.00%	0.64
RhythmNet(Fine-tuned)	0.43	3.97	3.99	4.06%	0.87

face videos. This is understandable because a single NIR channel may not convey as sufficient HR rhythm information as the three-channel RGB videos. In addition, we notice the SeetaFace face detector is trained using RGB images, and may fail to detect some faces in the NIR face images.

C. Cross-Database Testing

In order to validate the generalization ability of our RhythmNet, we also perform cross-database evaluations using our VIPL-HR database and two widely used HR estimation databases, i.e., MAHNOB-HCI, and MMSE-HR databases. Specifically, we train our RhythmNet on the color videos of VIPL-HR database and directly test it on the MAHNOB-HCI and MMSE-HR databases. We also fine-tune the VIPL-HR pre-trained model on MAHNOB-HCI and MMSE-HR following a three-fold subject-independent cross-validation testing protocol to see whether per-database fine-tuning could improve the HR estimation performance or not. All the results are provided in Table VII and Table VIII. The baseline methods we use for comparisons are Poh2010 [1], Poh2011 [2], Balakrishnan2013 [3], Li2014 [5], Haan2013 [4] and Tulyakov2016 [6], and we directly use the results of these methods provided in [5] and [6].

From the results, we can see that the proposed method could achieve promising results even when we directly test our VIPL-HR pre-trained model on these two databases, i.e., a RMSE of 8.28 bpm on the MAHNOB-HCI database, and a RMSE of 7.33 bpm on the MMSE-HR database. The error rates are further reduced to 3.99 bpm and 5.49 bpm on

TABLE VIII
THE HR ESTIMATION RESULTS BY THE PROPOSED APPROACH AND SEVERAL STATE-OF-THE-ART METHODS ON THE MMSE-HR DATABASE

Method	Mean (bpm)	Std (bpm)	RMSE (bpm)	MER	r
Li2014 [5]	11.56	20.02	19.95	14.64%	0.38
Haan2013 [4]	9.41	14.08	13.97	12.22%	0.55
Tulyakov2016 [6]	7.61	12.24	11.37	10.84%	0.71
RhythmNet(WithinDB)	-0.85	4.99	5.03	3.67%	0.86
RhythmNet(CrossDB)	-2.33	6.98	7.33	3.62%	0.78
RhythmNet(Fine-tuned)	-0.88	5.45	5.49	3.58%	0.84

MAHNOB-HCI and MMSE-HR, respectively when we fine-tune the pre-trained model on MAHNOB-HCI and MMSE-HR databases. Both results by the proposed approach are much better than previous methods. These results indicate that the variations of illumination, movement, and acquisition sensors covered in the VIPL-HR database are helpful for learning a deep HR estimator which has good generalization ability to unseen scenarios. In addition, the proposed RhythmNet is able to leverage the diverse image acquisition conditions in VIPL-HR to learn a robust HR estimator.

D. Key Components Analysis

In this section, we further analyze the effectiveness of our approach from multiple aspects, covering spatial-temporal map, color space selection, temporal modeling via GRU and computational cost. All the experiments are conducted on the VIPL-HR database under a subject-exclusive five-fold cross-validation protocol.

1) *Spatial-Temporal Map*: From the results in Table V, we can see that under the same experimental settings, I3D [38] trained from raw RGB face videos achieves a RMSE of 15.9 bpm, which is much worse than our RhythmNet with spatial-temporal maps as representations (RMSE = 8.14 bpm). These results indicate that it is very difficult for deep neural networks to directly learn informative features from the raw RGB videos. We further visualize the feature maps of I3D for the input video clips using the visualization method proposed in [39]. The visualization map is calculated based on the average of all the feature maps of all video clips, i.e., 16427 video clips from 476 videos in one fold of test. The map is shown in Fig. 15(a). The visualization map of our RhythmNet is given in Fig. 15(b). Since our RhythmNet takes the spatial-temporal maps as input, we calculated the mean responses of each rows (spatial dimension) and then reshaped the responses to a visualization map. From the visualization maps, we can see that our RhythmNet focuses on the informative face regions, and thus gives more accurate estimation results.

2) *Color Space Selection*: As discussed in [40], head movement effects much more on the intensity than the chromaticity of the image since the chromaticity reflects the intrinsic optical properties of hemoglobin in blood. Therefore, choosing a color space separating chromaticity from intensity is helpful to reduce the artifacts and thus benefit the training procedure.

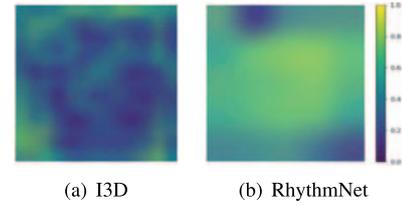


Fig. 15. The mean visualization maps of all the video clips in the test set of one fold for (a) I3D model [38] and (b) our RhythmNet.

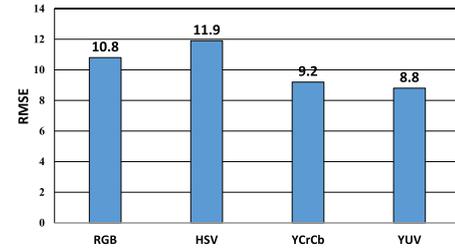


Fig. 16. The HR estimation errors (in RMSE) by RhythmNet using spatial-temporal maps generated from different color spaces (RGB, HSV, YCrCb, and YUV).

We considered three commonly used color space transformations, i.e., HSV, YCrCb, and YUV, as well as the RGB color space. For each experiment, the color space used for testing is the same as the color space used for training. The results by RhythmNet using these color spaces are given in Fig. 16.

From the results, we can see that different color spaces lead to very different HR estimation accuracies, and using YUV color space achieves the best result with a RMSE of 8.8 bpm. The results indicate that the YUV color space is more suitable in representing the HR signals, and helpful for capturing more information about heart rhythm. This is the reason that we choose to use the YUV color space in all the experiments.

3) *Temporal Modeling*: For HR measurement, we use GRU to model the relationship between adjacent measurements from succeeding video clips. We study the importance of this part by removing it from our RhythmNet and direct compute the average HR of estimations for all the clips. The results by RhythmNet with and without using the GRU are given in Table V. From the results, we can see that removing the GRU module leads to increased RMSE from 8.11 bpm to 8.94 bpm, which suggests the usefulness of the temporal modelling module in our RhythmNet.

4) *Computational Cost*: The model size of our HR estimator is about 42 MB, and the inference time is about 8ms on a Titan 1080Ti GPU when using a sliding window with 300 frames. These results indicate that the proposed approach is efficient, enabling real-time HR measurement.

E. Further Analysis

While the proposed RhythmNet achieves promising results on multiple databases, robust remote HR estimation under less-constrained scenarios faces a number of challenges. We provide additional analysis in terms of the challenging factors, i.e., video compression, illumination, head movement.

1) *Video Compression*: As shown in Section III-B, when using MJPG to compress the video data, the state-of-the-art

TABLE IX
THE HR ESTIMATION RESULTS BY OUR RHYTHMNET USING COMPRESSED AND UNCOMPRESSED RGB FACE VIDEOS OF VIPL-HR

Datatype	Mean (bpm)	Std (bpm)	MAE (bpm)	RMSE (bpm)	MER	r
Compressed	0.34	8.76	5.94	8.77	7.54%	0.58
Uncompressed	0.08	8.69	5.81	8.69	7.24%	0.59

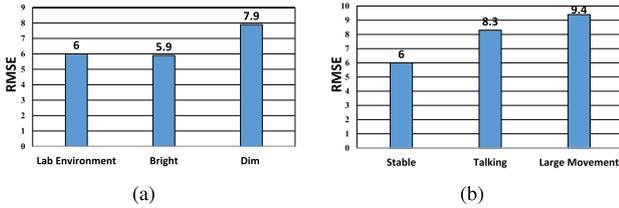


Fig. 17. The HR estimation errors (in RMSE) by RhythmNet under (a) different illumination conditions and (b) different head movement conditions.

method [4] shows minor performance degradation. We are also interested to see how such a compression can influence the proposed RhythmNet. As shown in Table IX, the HR estimation errors by our RhythmNet before and after compression are also the same. This suggests that the proposed approach is also very robust to the video compression by MJPG. From the result, we can see that the estimated results based on compressed data are very close to the results based on uncompressed data, which indicates that the compressed version VIPL-HR database is able to maintain the HR signals. All the experiments in this paper are based on the compressed version of VIPL-HR database.

2) *Illumination*: We study the influence of illumination variations on our RhythmNet by using the RGB face videos of Scenario 1, 4, and 5 in Table IV from the VIPL-HR dataset. The RMSE are reported in Fig. 17(a). From the results, we can see that under a bright lighting condition, i.e., with the lamp placed in front of the subject, our RhythmNet achieves a lower RMSE of 5.90 bpm comparing with the result with lab environment illumination. At the same time, under dim light conditions, the estimation error raises to 7.94 bpm. These results indicate that a good light condition is important for improving the HR estimation accuracy. However, even under dim light condition, the HR estimation error by our RhythmNet is lower than 8 bpm, which an encouraging result for practical application usage.

We also study the influence of RGB and NIR videos for HR estimation under dark scenario by user tests with 6 subjects. As shown in Fig. 18(a), we have recorded 8 RGB videos and 8 NIR videos for testing. The RhythmNet used for testing remains the one trained on the VIPL-HR. The HR estimation results using RGB and NIR videos are given in Fig. 18(b). As expected, we can see that under dark environment, RhythmNet gets an increased HR estimation error than under bright environment (from 8.14 bpm in RMSE to 17.7 bpm in RMSE on VIPL-HR). Using NIR videos for HR estimation can reduce the HR estimation error to $RMSE = 12$ bpm under the same dark environment. These

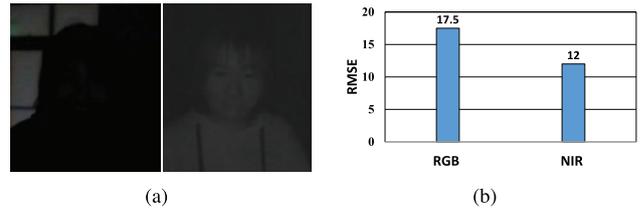


Fig. 18. (a) Examples of a RGB (left) and a NIR (right) face videos captured in dark scenario. (b) The corresponding HR estimation errors (in RMSE) by RhythmNet trained on RGB and NIR videos.

results indicate that NIR face videos can be a good choice for HR measurement in dark environment.

3) *Head Movement*: In order to analyse the influence of head movements, we compute the HR estimation errors for Situation 1, 2, and 3 in Table IV from the VIPL-HR dataset. The results are shown in Fig. 17(b). From the results, we can see that all kinds of head movements could significantly influence the estimation accuracy. The RMSE raises to 8.28 bpm when the subject is talking, and 9.40 bpm when the subject performs large scale head movements. These results indicate that HR estimation under different head movements is still a challenging task, and should be considered when building a robust HR estimator.

VI. CONCLUSION AND FUTURE WORK

While remote HR estimation from a face video has wide applications, it is a challenging problem in less-constrained scenarios due to variations of head movement, illumination, and sensor diversity. In this paper, we propose an end-to-end learning approach (RhythmNet) for robust heart rate estimation from face videos. Our RhythmNet uses a novel spatial-temporal map represent the HR signals in the videos, which allows efficient modelling via CNN. We also take into account the temporal relationship of adjacent HR estimates to perform continuous HR measurement. At the same time, we also build a multi-modality database (VIPL-HR) for studying HR estimation under less-constrained scenarios covering variations of illumination, head movement, and sensor diversity. The proposed approach achieves promising HR estimation accuracies in both within-database and cross-database testing scenarios on the MAHNOB-HCI, MMSE-HR, and VIPL-HR databases. In our future work, we would like to investigate the effectiveness of the proposed approach for the other rPPG-based physiological status measurement tasks, such as the breath rate and blood pressure measurement from videos. We would also like to investigate robust HR estimation approaches leveraging distribution learning [41] or multi-task learning [42], [43].

ACKNOWLEDGMENT

The authors would like to thank the editors and anonymous reviewers for their constructive comments and suggestions.

REFERENCES

- [1] M.-Z. Poh, D. J. McDuff, and R. W. Picard, "Non-contact, automated cardiac pulse measurements using video imaging and blind source separation," *Opt. Express*, vol. 18, no. 10, pp. 10762–10774, 2010.

- [2] M.-Z. Poh, D. J. McDuff, and R. W. Picard, "Advancements in noncontact, multiparameter physiological measurements using a webcam," *IEEE Trans. Biomed. Eng.*, vol. 58, no. 1, pp. 7–11, Jan. 2011.
- [3] G. Balakrishnan, F. Durand, and J. Guttag, "Detecting pulse from head motions in video," in *Proc. IEEE CVPR*, Jun. 2013, pp. 3430–3437.
- [4] G. de Haan and V. Jeanne, "Robust pulse rate from chrominance-based rPPG," *IEEE Trans. Biomed. Eng.*, vol. 60, no. 10, pp. 2878–2886, Oct. 2013.
- [5] X. Li, J. Chen, G. Zhao, and M. Pietikäinen, "Remote heart rate measurement from face videos under realistic situations," in *Proc. IEEE CVPR*, Jun. 2014, pp. 4264–4271.
- [6] S. Tulyakov, X. Alameda-Pineda, E. Ricci, L. Yin, J. F. Cohn, and N. Sebe, "Self-adaptive matrix completion for heart rate estimation from face videos under realistic conditions," in *Proc. IEEE CVPR*, Jun. 2016, pp. 2396–2404.
- [7] W. Wang, A. C. den Brinker, S. Stuijk, and G. de Haan, "Algorithmic principles of remote PPG," *IEEE Trans. Biomed. Eng.*, vol. 64, no. 7, pp. 1479–1491, Jul. 2017.
- [8] M. Soleymani, J. Lichtenauer, T. Pun, and M. Pantic, "A multimodal database for affect recognition and implicit tagging," *IEEE Trans. Affect. Comput.*, vol. 3, no. 1, pp. 42–55, Jan. 2012.
- [9] G.-S. Hsu, A. Ambikapathi, and M.-S. Chen, "Deep learning with time-frequency representation for pulse estimation from facial videos," in *Proc. IEEE IJCB*, Oct. 2017, pp. 642–650.
- [10] R. Stricker, S. Müller, and H.-M. Gross, "Non-contact video-based pulse rate measurement on a mobile service robot," in *Proc. IEEE RO-MAN*, Aug. 2014, pp. 1056–1062.
- [11] R. Irani, K. Nasrollahi, and T. B. Moeslund, "Improved pulse detection from head motions using DCT," in *Proc. VISAPP*, vol. 3, Jan. 2014, pp. 118–124.
- [12] X. Li *et al.*, "The OBF database: A large face video database for remote physiological signal measurement and atrial fibrillation detection," in *Proc. IEEE FG*, May 2018, pp. 1–6.
- [13] X. Niu, H. Han, S. Shan, and X. Chen, "VIPL-HR: A multi-modal database for pulse estimation from less-constrained face video," in *Proc. ACCV*, 2018, pp. 562–576.
- [14] D. McDuff, S. Gontarek, and R. W. Picard, "Improvements in remote cardiopulmonary measurement using a five band digital camera," *IEEE Trans. Biomed. Eng.*, vol. 61, no. 10, pp. 2593–2601, Oct. 2014.
- [15] A. Lam and Y. Kuno, "Robust heart rate measurement from video using select random patches," in *Proc. IEEE ICCV*, Dec. 2015, pp. 3640–3648.
- [16] M. Lewandowska, J. Rumiński, T. Kocejko, and J. Nowak, "Measuring pulse rate with a webcam—A non-contact method for evaluating cardiac activity," in *Proc. FedCSIS*, Sep. 2011, pp. 405–410.
- [17] G. R. Tsouri, S. Kyal, S. A. Dianat, and L. K. Mestha, "Constrained independent component analysis approach to nonobtrusive pulse rate measurements," *J. Biomed. Opt.*, vol. 17, no. 7, 2012, Art. no. 077011.
- [18] W. Wang, S. Stuijk, and G. de Haan, "Exploiting spatial redundancy of image sensor for motion robust rPPG," *IEEE Trans. Biomed. Eng.*, vol. 62, no. 2, pp. 415–425, Feb. 2015.
- [19] L. Feng, L. M. Po, X. Xu, Y. Li, and R. Ma, "Motion-resistant remote imaging photoplethysmography based on the optical properties of skin," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 25, no. 5, pp. 879–891, May 2015.
- [20] G. De Haan and A. Van Leest, "Improved motion robustness of remote-PPG by using the blood volume pulse signature," *Physiol. Meas.*, vol. 35, no. 9, p. 1913, 2014.
- [21] M. Kumar, A. Veeraraghavan, and A. Sabharwal, "DistancePPG: Robust non-contact vital signs monitoring using a camera," *Biomed. Opt. Express*, vol. 6, no. 5, pp. 1565–1588, 2015.
- [22] W. J. Wang, A. C. den Brinker, S. Stuijk, and G. de Haan, "Amplitude-selective filtering for remote-PPG," *Biomed. Opt. Express*, vol. 8, no. 3, pp. 1965–1980, 2017.
- [23] X. Niu, H. Han, S. Shan, and X. Chen, "Continuous heart rate measurement from face: A remote rPPG approach with distribution learning," in *Proc. IEEE IJCB*, Oct. 2017, pp. 642–650.
- [24] W. Wang, S. Stuijk, and G. de Haan, "A novel algorithm for remote photoplethysmography: Spatial subspace rotation," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 9, pp. 1974–1984, Sep. 2016.
- [25] Y. Hsu, Y.-L. Lin, and W. Hsu, "Learning-based heart rate detection from remote photoplethysmography features," in *Proc. IEEE ICASSP*, May 2014, pp. 4433–4437.
- [26] W. Chen and D. McDuff, "DeepPhys: Video-based physiological measurement using convolutional attention networks," *Proc. ECCV*, 2018, pp. 349–365.
- [27] M. V. Gastel, S. Stuijk, and G. D. Haan, "Motion robust remote-PPG in infrared," *IEEE Trans. Biomed. Eng.*, vol. 62, no. 5, pp. 1425–1433, May 2015.
- [28] J. Chen *et al.*, "RealSense = real heart rate: Illumination invariant heart rate estimation from videos," in *Proc. IEEE IPTA*, Dec. 2016, pp. 1–6.
- [29] H.-Y. Wu, M. Rubinstein, E. Shih, J. Guttag, F. Durand, and W. Freeman, "Eulerian video magnification for revealing subtle changes in the world," *ACM Trans. Graph.*, vol. 31, no. 4, pp. 1–8, 2012.
- [30] W. Verkruysse, L. O. Svaasand, and J. S. Nelson, "Remote plethysmographic imaging using ambient light," *Opt. Express*, vol. 16, no. 26, pp. 21434–21445, 2008.
- [31] C. Yang, G. Cheung, and V. Stankovic, "Estimating heart rate and rhythm via 3D motion tracking in depth video," *IEEE Trans. Multimedia*, vol. 19, no. 7, pp. 1625–1636, Jul. 2017.
- [32] D. J. McDuff, E. B. Blackford, and J. R. Estepp, "The impact of video compression on remote cardiac pulse measurement using imaging photoplethysmography," in *Proc. IEEE FG*, May/Jun. 2017, pp. 63–70.
- [33] G. R. Tsouri and Z. Li, "On the benefits of alternative color spaces for noncontact heart rate measurements using standard red-green-blue cameras," *J. Biomed. Opt.*, vol. 20, no. 4, 2015, Art. no. 048002.
- [34] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE CVPR*, Jun. 2016, pp. 770–778.
- [35] K. Cho *et al.*, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," 2014, *arXiv:1406.1078*. [Online]. Available: <https://arxiv.org/abs/1406.1078>
- [36] X. Niu *et al.*, "Automatic engagement prediction with GAP feature," in *Proc. ICMI*, 2018, pp. 599–603.
- [37] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*. [Online]. Available: <https://arxiv.org/abs/1412.6980>
- [38] J. Carreira and A. Zisserman, "Quo vadis, action recognition? A new model and the kinetics dataset," in *Proc. IEEE CVPR*, Jul. 2017, pp. 6299–6308.
- [39] Z. Yu, X. Li, and G. Zhao, "Remote photoplethysmograph signal measurement from facial videos using Spatio-temporal networks," in *Proc. BMVC*, 2019, pp. 1–12. [Online]. Available: <https://bmvc2019.org/wp-content/uploads/papers/0186-paper.pdf>
- [40] Y. Yang, C. Liu, H. Yu, D. Shao, F. Tsow, and N. Tao, "Motion robust remote photoplethysmography in CIE Lab color space," *J. Biomed. Opt.*, vol. 21, no. 11, 2016, Art. no. 117001.
- [41] H. Pan, H. Han, S. Shan, and X. Chen, "Mean-variance loss for deep age estimation from a face," in *Proc. IEEE/CVF CVPR*, Jun. 2018, pp. 5285–5294.
- [42] F. Wang, H. Han, S. Shan, and X. Chen, "Deep multi-task learning for joint prediction of heterogeneous face attributes," in *Proc. IEEE FG*, May/Jun. 2017, pp. 173–179.
- [43] J. Cui, H. Zhang, H. Han, S. Shan, and X. Chen, "Improving 2D face recognition via discriminative face depth estimation," in *Proc. IEEE ICB*, Feb. 2018, pp. 140–147.



Xuesong Niu received the B.E. degree from Nankai University. He is currently pursuing the Ph.D. degree with the Institute of Computing Technology (ICT), Chinese Academy of Sciences (CAS), Beijing, China. His research interests include computer vision, machine learning, and affective computing.



Shiguang Shan (M'04–SM'15) received the Ph.D. degree in computer science from the Institute of Computing Technology (ICT), Chinese Academy of Sciences (CAS), Beijing, China, in 2004, where he has been a Full Professor since 2010 and currently the Deputy Director of the CAS Key Laboratory of Intelligent Information Processing. He is also a member of the CAS Center for Excellence in Brain Science and Intelligence Technology. He has published over 300 articles, with totally over 15 000 Google scholar citations. His research interests cover computer vision, pattern recognition, and machine learning. He served as an Area Chair for many international conferences, including ICCV11, ICASSP14, ICPR12/14/19, ACCV12/16/18, FG13/18/20, BTAS18, and CVPR19/20. He was/is an Associate Editor of several journals, including IEEE T-IP, *Neurocomputing*, CVIU, and PRL. He was a recipient of the China's State Natural Science Award in 2015 and the China's State S&T Progress Award in 2005 for his research work.



Hu Han received the B.S. degree in computer science from Shandong University and the Ph.D. degree in computer science from the Institute of Computing Technology (ICT), Chinese Academy of Sciences (CAS), in 2005 and 2011, respectively. He was a Research Associate with the PRIP Lab, Department of Computer Science and Engineering, Michigan State University, and a Visiting Researcher with Google, Mountain View. He joined the faculty at ICT, CAS, in 2015. He is currently an Associate Professor with the Institute of Computing Technology (ICT), Chinese Academy of Sciences (CAS). He has authored or coauthored over 50 articles in refereed journals and conferences, including IEEE Trans. PAMI/IFS, CVPR, ECCV, NeurIPS, and MICCAI. His research interests include computer vision, pattern recognition, and image processing, with applications to biometrics and medical image analysis. He was a recipient of the IEEE FG2019 Best Poster Presentation Award and CCBP 2016/2018 Best Student/Poster Awards.



Xilin Chen is currently a Professor with the Institute of Computing Technology, Chinese Academy of Sciences (CAS). He has authored one book and over 300 articles in refereed journals and proceedings in the areas of computer vision, pattern recognition, image processing, and multimodal interfaces. He is a fellow of IAPR and CCF. He served as an Organizing Committee Member for many conferences, including a General Co-Chair of FG13/FG18 and a Program Co-Chair of ICMI 2010. He is/was an Area Chair of CVPR 2017/2019/2020, and ICCV 2019. He is currently an Associate Editor of the IEEE TRANSACTIONS ON MULTIMEDIA, a Senior Editor of the *Journal of Visual Communication and Image Representation*, a Lead Editor of the *Journal of Computer Science and Technology*, and an Associate Editor-in-Chief of the *Chinese Journal of Computers* and *Chinese Journal of Pattern Recognition and Artificial Intelligence*.