

Neuron Structure Modeling for Generalizable Remote Physiological Measurement

Hao LU^{1,2}, Zitong YU³, Xuesong NIU⁴, Yingcong CHEN^{1,2,*}

¹The Hong Kong University of Science & Technology (Guangzhou),

²The Hong Kong University of Science & Technology, ³Great Bay University, ⁴Kuaishou Technology
hlu585@connect.hkust-gz.edu.cn, zitong.yu@ieee.org,
nxsedson@gmail.com, yingcongchen@ust.hk

Abstract

Remote photoplethysmography (rPPG) technology has drawn increasing attention in recent years. It can extract Blood Volume Pulse (BVP) from facial videos, making many applications like health monitoring and emotional analysis more accessible. However, as the BVP signal is easily affected by environmental changes, existing methods struggle to generalize well for unseen domains. In this paper, we systematically address the domain shift problem in the rPPG measurement task. We show that most domain generalization methods do not work well in this problem, as domain labels are ambiguous in complicated environmental changes. In light of this, we propose a domain-label-free approach called NEuron STructure modeling (NEST). NEST improves the generalization capacity by maximizing the coverage of feature space during training, which reduces the chance for under-optimized feature activation during inference. Besides, NEST can also enrich and enhance domain invariant features across multi-domain. We create and benchmark a large-scale domain generalization protocol for the rPPG measurement task. Extensive experiments show that our approach outperforms the state-of-the-art methods on both cross-dataset and intra-dataset settings. The codes are available at <https://github.com/LuPaoPao/NEST>.

1. Introduction

Physiological signals such as heart rate (HR), and heart rate variability (HRV), respiration frequency (RF) are important body indicators that serve not only as vital signs but also track the level of sympathetic activation [17, 33, 54]. Traditional physiological measurements, such as electrocardiograms, heart rate bands, and finger clip devices, have high accuracy. However, they are costly, intrusive, and uncomfortable to wear for a long time. Remote photoplethysmography (rPPG) can extract blood volume pulse (BVP)

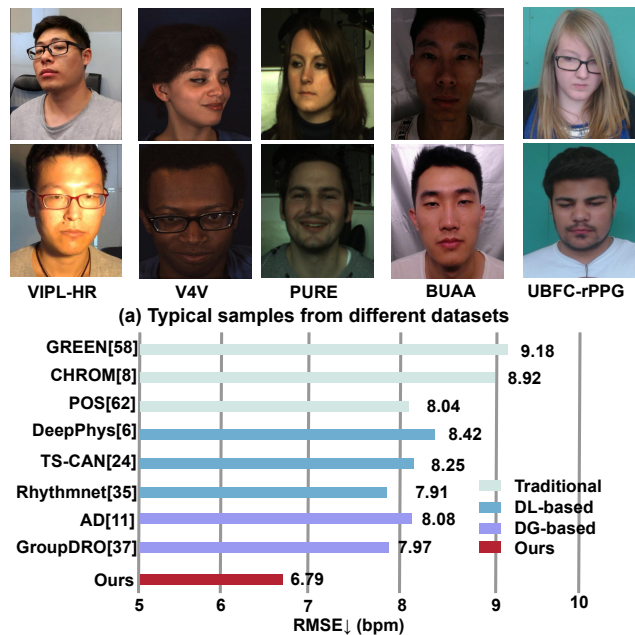


Figure 1. (a) Typical samples from different publicly rPPG datasets: VIPL-HR [36], V4V [47], UBFC-rPPG [1], BUAA [68], PURE [52]. (b) The performance of different methods on DG protocol (test on the UBFC-rPPG dataset with training on the VIPL, V4V, PURE, and BUAA).

from face video, which analyzes the periodic changes of the light absorption of the skin caused by heartbeats. Then various physiological indicators (such as HR and HRV) can be calculated based on BVP signals [37, 69, 70]. With non-intrusion and convenience, the rPPG-based physiological measurement method can only use an ordinary camera to monitor physiological indicators and gradually become a research hotspot in the computer vision field [7, 28, 32, 36, 44, 49].

Traditional rPPG measurement methods include signal blind decomposition [21, 34, 42] and color space transformation [9, 59, 63]. These approaches rely on heartbeat-related

*Corresponding author.

statistical information, only applicable in constrained environments. In recent years, deep learning (DL) based approaches [6, 19, 25, 28, 36, 44, 50, 69, 70] have shown their great potentials in rPPG measurement. By learning dedicated rPPG feature representation, these methods achieve promising performance in much more complicated environments [1, 36, 47, 52, 68].

However, deep learning methods suffer from significant performance degradation when applied in real-world scenarios. This is because most training data are captured in lab environments with limited environmental variations. With domain shifts (e.g., different illumination, camera parameters, motions, etc.), these models may struggle to generalize for the unseen testing domain. To validate this, we conduct a cross-dataset evaluation shown in Fig. 1(b). As shown, all DL-based methods do not work well in this evaluation. Furthermore, it is worth noting that DeepPhys [6] and TS-CAN [25] even perform inferior to traditional approach POS [63].

To improve the performance on the unseen domains, one common practice is to incorporate Domain Generalization (DG) approaches, e.g., encouraging intermediate features to be domain-invariant [12, 18, 30, 38, 57, 62]. However, as shown in Fig. 1 (b) (AD [12] and GroupDRO [38]), the improvements are still quite limited. One reason is that existing DG methods assume that domain labels can be clearly defined (e.g., by data source). Unfortunately, different background scenes, acquisition devices, or even identities could cause very different data distributions in rPPG measurement tasks, and such distribution discrepancy may exist within or cross datasets. In this case, explicitly defining domains is very difficult, and simply treating one dataset as one domain may lead to inferior performance. This is termed as *agnostic domain generalization*. Besides this, as physiological cues are usually much more subtle than the various noise, the model may overfit in the source domain with the limited training data.

In this paper, we propose the NEural SStructure modeling (NEST), a principled solution to the abovementioned problems. The main idea of NEST is to narrow the under-optimized and redundant feature space, align domain invariant features, and enrich discriminative features. Our intuition is as follows. Neural structure refers to the channel activation degree in each convolution layer, which reveals the discriminative feature combination for the specific sample. As the limited variation in a certain restricted domain, there are some spaces that the model seldomly is optimized. Out-of-distribution (OOD) samples may cause abnormal activation in these spaces, which may lead to performance degeneration. Therefore, we regularize the neural structure to encourage the model to be more well-conditioned to avoid abnormal activation caused by OOD samples. Specifically, we propose the NEural SStructure Coverage Maxi-

mization (NEST-CM) that encourages all neural spaces to be optimized during training, reducing the chance of abnormal activation during testing. Secondly, we propose the NEural SStructure Targeted Alignment (NEST-TA) that encourage network suppresses domain variant feature by comparing the samples with the similar physiological information. Thirdly, we propose the NEural SStructure Diversity Maximization (NEST-DM) to enrich discriminative features against unseen noise. It should be noted that our approach does not rely on domain labels, which is more applicable in the rPPG measurement task. To summarize, our contributions are listed as follows:

1. We are the first to study the domain shift problem in rPPG measurement, which introduces a new challenge, agnostic domain generalization.
2. We propose the NEural SStructure modeling to alleviate domain shift, which is performed by narrowing the under-optimized feature space, and enhancing and enriching domain invariant features.
3. We establish a large-scale domain generalization (DG) benchmark for rPPG measurement, which is the first DG protocol in this task. Extensive experiments in this dataset show the superiority of our approach.

2. Related work

2.1. Remote Physiological Measurement

Remote physiological measurement aims to estimate HR and HRV values by analyzing the chrominance changes in the skin. Traditional methods mainly perform color space transformations or signal decomposition to obtain the BVP signal with high SNR [9, 10, 42, 64]. These methods are designed manually for certain lighting scenes with slight motion and can not generalize to the complex environment in the wild. With powerful modeling capacity, the deep learning (DL) model has been successfully used in remote physiological estimation [6, 13, 35, 36, 51, 67, 69]. Among these methods, various model structures are tailored to extract BVP signals from face videos by aggregating spatial and temporal information [6, 37, 51, 69, 70]. Besides, many hand-designed feature maps are used to reduce the computational burden of the network by removing the irrelevant background [13, 28, 36, 45, 50]. Some work utilizes the distribution of the target domain to improve the model performance [19, 26], which is not practical without touching the unseen domain in advance. Differently, we study the model generalization on domain generalization protocol for rPPG measurement.

2.2. Domain generalization

Domain generalization (DG), i.e., out-of-distribution generalization, aims to generalize to the unseen domain by training a model on multiple or single source domains [62]. The main solution can be categorized into data manipulation, representation learning, and meta-learning. Data

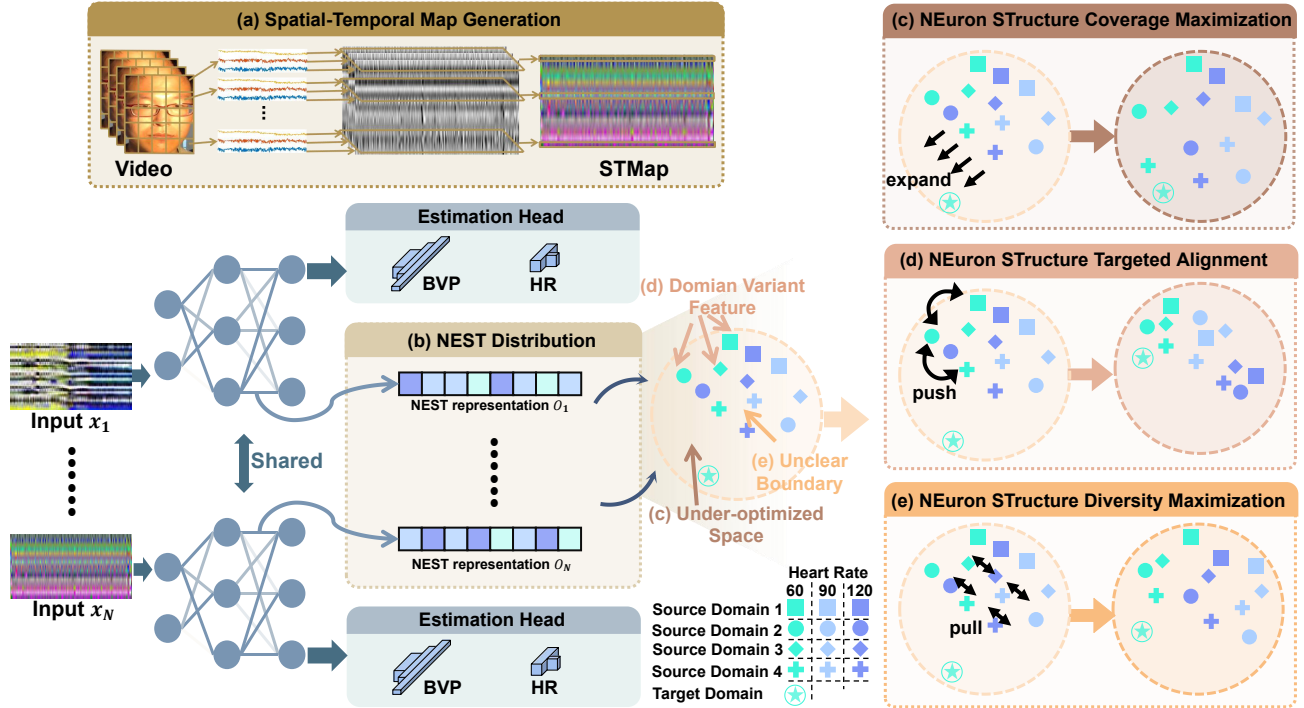


Figure 2. An overview of the proposed method. (a) Spatial-temporal map (STMap) is extracted from face video [36]. (b) NEST representation is calculated from the samples of the different domains. (c) The under-optimized space in NEST presentation distribution is narrowed by NEST-CM. (d) NEST-TA align the feature of multi domain at instance-level. (e) The distance of different samples are pulled away by NEST-DM.

manipulation is an important technique to improve the generalization of the model, which can remove irrelevant style information by data generation and data augmentation [48, 71]. The main goal of representation learning is to learn domain-invariant features by alleviating covariate shift across domain [12, 18, 38, 57]. By simulating domain shift on the virtual dataset, meta-learning seeks transferable and shareable knowledge from multiple domain [23, 30]. Unlike common tasks such as classification, segmentation, and detection [15, 20, 22, 24, 40], rPPG measurement tasks do not have explicit boundaries between different domains, i.e., agnostic domain generalization. Recently, some single source domain generalization (SSDG) methods focus on the worst-case formulation without using the domain label [43, 57, 66]. Different from all the methods above, we tackle the DG problem of regression task in view of neuron structure modeling.

3. Method

3.1. Overall Framework

In this work, we propose the NEuron SStructure modeling (NSET), without using domain label, to alleviate the performance degradation caused by domain shift as illustrated in Fig. 2. Specifically, we first compressed the face video into the spatial-temporal map (STMap) as the input of

CNN refer to [36] as shown in Fig. 2 (a). Sequentially, we define an Activated NEuron SStructure (NEST) representation that describes the activated degree of different channels of the network as shown in Fig. 2 (b). To alleviate the domain shift, we model a unified NEST distribution by three regularizations: (1) NEuron SStructure Coverage Maximization (NEST-CM) is used to narrow under-optimized space that is redundant and unstable as shown in Fig. 2 (c); (2) NEuron SStructure Targeted Alignment (NEST-TA) is introduced to align multi-domain feature distribution at instance-level as shown in Fig. 2 (d). (3) NEuron SStructure Diversity Maximization (NEST-DM) loss to force the network to explore more features with discriminant ability as shown in Fig. 2 (e).

3.2. NEuron SStructure Modeling

NEuron SStructure Representation. For the rPPG measurement task, the convolution kernels are adaptively activated to extract physiological features against various noise distribution [27, 56, 65]. However, the facial physiological signals are feeble and can be easily affected by various unseen noises, which make the filters activated abnormally. Therefore, the NEuron SStructure (NEST) representation is defined to explore how domain shift affects the model robustness in view of activated feature distribution. Specifically, we firstly average the feature map

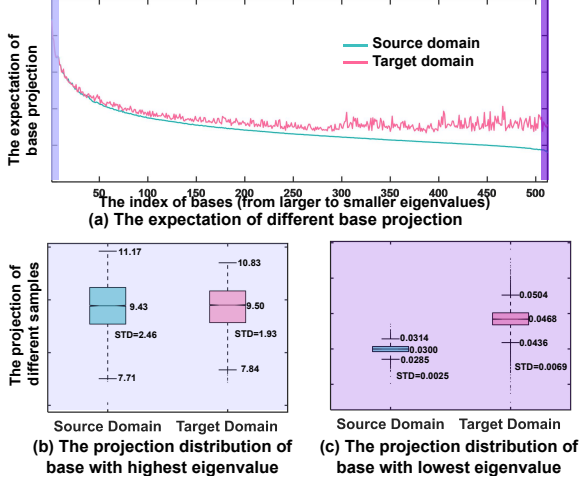


Figure 3. (a) The expectation of different base projection. (b) The distribution of base with highest eigenvalue. (c) The distribution of base with lowest eigenvalue.

$m_j \in R^{W \times H \times C_j}$ over the spatial dimensions in the end of j -th layer to get the preliminary statistical presentation $\bar{o}_j \in R^{C_j}$ of the neuron structure, where C_j is the channel number of the j -th layer. Following that, the preliminary presentation \bar{o}_j is fed into the max-min normalized to get the NEST representation $o_j \in R_j^C$ of the j -th layer. The full NEST representation of i -th sample x_i is defined as $O_i = \{o_{i,1}, o_{i,2}, \dots, o_{i,U}\} \in R^{(C_1+C_2+\dots+C_U)}$, where U is the number of selected layer.

NEuron STructure Coverage Maximization. Existing works show that the OOD samples may activate abnormal neuron structures (feature combination) that seldom be trained with the source data so that the model’s performance degrades drastically [31, 39, 57]. To investigate this phenomenon in rPPG measurement task, we analyze the optimized degree of different feature combinations.

Specifically, we decompose the NEST representation of the last layer into orthogonal bases with Singular Value Decomposition (SVD). Here, small eigenvalues indicate that training samples cover a small space in these corresponding bases. Thus, unseen target samples are more likely to lay out of the well-optimized space, causing a larger domain shift. To validate this, we plot the means of source and target samples on these bases in Fig. 3(a). As shown, the mean differences are larger for bases of smaller eigenvalues. To make it clearer, We further visualize the data distribution on the first and last bases in Fig. 3 (b) and (c). As shown, the distribution difference in the last base is much more significant than that of the first one. All these suggest that for the NEST representation, components of small eigenvalues are less optimized during training, and domain shifts are more likely to appear there.

In sight of this, we propose the NEuron STructure Coverage Maximization (NEST-CM) regularization that improves




Figure 4. The detail of SVD for the j -th layer in CNN.

the eigenvalues of the last few bases. This could narrow the under-optimized space as shown in Fig. 2 (c). For a training batch with N samples, NEST representation set of j -th layer $B_j = \{o_{1,j}; o_{2,j}; \dots; o_{N,j}\} \in R^{N \times C_j}$ is decomposed to $S_j \cdot V_j \cdot D_j$ by SVD method as shown in Fig. 4, i.e., $B_j = S_j \cdot V_j \cdot D_j$, where $V_j \in R^{N \times C_j}$ and $N > C_j$. The bases with small eigenvalue values are rare and under-optimized during the training stage. To reduce the occurrence probability of under-optimized NEST caused by domain shift, we stimulate under-optimized bases by maximizing the corresponding eigenvalues:

$$\mathcal{L}_{CM} = -\mathbb{E}_{j \in U} \frac{\sum_{i=1}^{C'} \lambda'_{i,j}}{\sum_{i=1}^C \lambda_{i,j}}, \lambda'_{i,j} < t, \quad (1)$$

where $\lambda_{i,j}$ represent the i -th normalized eigenvalue in j -th layer’ diagonal matrix V_j , and $\lambda'_{i,j}$ is the eigenvalue less than the threshold t . Considering different channel numbers in each layer, we design an adaptive threshold t to select the relatively under-optimized bases, i.e., $t = \frac{\rho}{C_j}$, where ρ is a trade-off parameter. Besides, C'_j is the number of eigenvalues of j -th layer under the threshold, and U is the number of layers.

NEuron STructure Targeted Alignment. Besides the NEST-CM introduced above, we further improve the generalization ability by aligning the NEST representation of different domains. Domain invariant learning has demonstrated its effectiveness in DG problems [12, 18, 30, 38, 57, 62]. Some recent research [11, 16, 29, 60] further shows that aligning domain features within the same labels could lead to better performance. However, labels in rPPG are continuous, and it is hard to find samples with identical labels in a training batch.

To cope with this, we propose a novel regularization term called NEuron STructure Alignment (NEST-TA). NEST-TA builds a synthesized prototype O_i^p with the same physiological information as the given targeted sample O_i . Then, the feature within the same labels can be aligned across domains like [11, 16, 29, 60]. Specifically, we select the top K samples closest to the targeted ground truth y_i , i.e., $[O'_{i,1}, \dots, O'_{i,k}, \dots, O'_{i,K}]$. The NEST center of the selected samples \bar{O}'_i is calculated by averaging the NEST presentation of selected samples. Then we calculate the direction vectors between the NEST presentation center of the selected samples \bar{O}'_i and the NEST presentation of each sample $O'_{i,k}$. Following that, the NEST center of the selected samples \bar{O}'_i is calibrated to the synthesized prototype O_i^p ,

which is performed by using the Gaussian distribution function to fuse these direction vectors:

$$O_i^p = \overline{O}_i + \sum_{k=1}^K W_{i,k} * \frac{y_i - \overline{y}_i'}{y'_{i,k} - \overline{y}_i'} (O'_{i,k} - \overline{O}_i), \quad (2)$$

$$W_{i,k} = \frac{G(y_{i,k} - \overline{y}_i', \sigma)}{\sum_{k=1}^K G(y_{i,k} - \overline{y}_i', \sigma)},$$

where $G(-, \sigma)$ is the Gaussian distribution with the mean of zero and the standard deviation of σ , i.e., $G(y_{i,k} - \overline{y}_i', \sigma) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(y_{i,k} - \overline{y}_i')^2}{2\sigma^2}}$, where σ is a hyper-parameter. To align the feature distribution across different domains, NEST-TA uses the cosine similarity $Cos(-, -)$ to reduce the variant features between targeted samples O_i and prototype sample O_i^p :

$$\mathcal{L}_{TA} = \mathbb{E}_{i \in N} Cos(O_i, O_i^p). \quad (3)$$

NEuron STructure Diversity Maximization. Existing works elaborate that it is effective to explore more features with discriminant information to improve the generalization [4,30,41,46]. Motivated by this, NEuron STructure Diversity Maximization (NEST-DM) is introduced to increase the distinction between the NEST presentation of different samples as shown in Fig. 2 (e). By doing this, the network is forced to explore more features with physiological information in a contrastive learning way, which can further improve model discrimination. NEST-DM \mathcal{L}_{DM} can be represented as:

$$\mathcal{L}_{DM} = -\mathbb{E}_{i \in N} \log\left(\frac{e^{Cos(O_i, \widehat{O}_i)/\tau}}{\sum_{j=1(i \neq j)}^N e^{Cos(O_i, \widehat{O}_j)/\tau}}\right), \quad (4)$$

where \widehat{O}_i is the NEST representation of augmented sample from O_i , and N and τ are the batch size and temperature parameter. The augmentation strategies for STMap will be elaborated in section 4.2.

Model Training. For the rPPG measurement task, we chose two popular losses in this field. Specifically, L1 loss \mathcal{L}_{L1} and negative Pearson’s correlation coefficient \mathcal{L}_p are used to constrain the predicted heart rate value and the predicted BVP signal respectively [14, 28, 37, 69]. These two losses are also used for augmented data. With the violent confrontation at the beginning of training, \mathcal{L}_{L1} make network tends to predict the mean of the heart rate value distribution instead of learning the physiological feature like \mathcal{L}_p . Besides, \mathcal{L}_{NSM} , \mathcal{L}_{NEST-S} , and $\mathcal{L}_{NEST-DM}$ also have meaningless effects on under-trained networks at the beginning.

To solve the above-illustrated issues, we introduce the adaptation factor $\gamma = \frac{2}{1 + \exp(-10r)}$ and $r = \frac{interscurrent}{interstotal}$ to suppress meaningless constraints at the early training stage, which can accelerate model convergence and improve training stability [12, 15]. The final loss $\mathcal{L}_{overall}$ can be formulated as follows:

$$\mathcal{L}_{overall} = k_1 \mathcal{L}_p + \gamma(k_2 \mathcal{L}_{L1} + k_3 \mathcal{L}_{CM} + k_4 \mathcal{L}_{TA} + k_5 \mathcal{L}_{DM}), \quad (5)$$

where $k_1 - k_5$ are five trade-off parameters.

4. Experiments

4.1. Datasets

We elaborately select five datasets (including different motion, camera, and lighting conditions) to curate and benchmark a large-scale DG protocol:

VIPL-HR [36] have nine scenarios, three RGB cameras, different illumination conditions, and different levels of movement. It is worth mentioning that the BVP signal and video of the dataset do not match in the time dimension, so \mathcal{L}_p is not used for this dataset. We used the output of HR head as the evaluation result. Besides, we normalize STMap to 30 fps by cubic spline interpolation to solve the problem of the unstable frame rate of video [28, 50, 70].

V4V [47] is designed to collect data with the drastic changes of physiological indicators by simulating ten tasks such as a funny joke, 911 emergency call, and odor experience. It is worth mentioning that There are only heart rate labels in the dataset and no BVP signal labels, so \mathcal{L}_p is not used for this dataset. We used the output of HR head as the evaluation result.

BUAA [68] is proposed to evaluate the performance of the algorithm against various illumination. We only use data with illumination greater than or equal to 10 lux because underexposed images require special algorithms that are not considered in this article. We used the output of BVP head as the evaluation for BUAA, PURE, and UBFC-rPPG datasets.

PURE [52] contains 60 RGB videos from 10 subjects with six different activities, specifically, sitting still, talking, and four rotating and moving head variations. The BVP signals are down-sampled from 60 to 30 fps with cubic spline interpolation to align with the videos.

UBFC-rPPG [1] containing 42 face videos with sunlight and indoor illumination, the ground-truth BVP signals, and HR values were collected by CMS50E.

4.2. Implementation Details

Training Details. Our proposed method is performed under the Pytorch framework. For the generation of STMap from Video, we use the FAN¹ to detect the 2D landmarks of faces [2], and other steps completely follow [36]. STMap in each dataset is sampled with a time window of 256 with step 10. Following [55, 61], we also design spatial and temporal augmentation for STMap, i.e., shuffle each row of STMap and slide the time window. In addition, color jitter and blur

¹<https://github.com/ladrianb/face-alignment>

Method	UBFC			PURE			BUAA			VIPL			V4V		
	MAE↓	RMSE↓	r↑	MAE↓	RMSE↓	r↑	MAE↓	RMSE↓	r↑	MAE↓	RMSE↓	r↑	MAE↓	RMSE↓	r↑
GREEN [59]	8.02	9.18	0.36	10.32	14.27	0.52	5.82	7.99	0.56	12.18	18.23	0.25	15.64	21.43	0.06
CHROM [9]	7.23	8.92	0.51	9.79	12.76	0.37	6.09	8.29	0.51	11.44	16.97	0.28	14.92	19.22	0.08
POS [63]	7.35	8.04	0.49	9.82	13.44	0.34	5.04	7.12	0.63	14.59	21.26	0.19	17.65	23.22	0.04
DeepPhys [6]	7.82	8.42	0.54	9.34	12.56	0.55	4.78	6.74	0.69	12.56	19.13	0.14	14.52	19.11	0.14
TS-CAN [25]	7.63	8.25	0.55	9.12	12.38	0.57	4.84	6.89	0.68	12.34	18.94	0.16	14.77	19.96	0.12
Rhythmnet * [36]	5.79	7.91	0.78	7.39	10.49	0.77	3.38	5.17	0.84	8.97	12.16	0.49	10.16	14.57	0.34
Dual-GAN * [28]	5.55	7.62	0.79	7.24	10.27	0.78	3.41	5.23	0.84	8.88	11.69	0.50	10.04	14.44	0.35
BVPNet * [7]	5.43	7.71	0.80	7.23	10.25	0.78	3.69	5.48	0.81	8.45	11.64	0.51	10.01	14.35	0.36
AD *+ [12]	5.92	8.08	0.76	7.42	10.61	0.73	3.49	5.49	0.82	8.41	11.71	0.53	10.47	14.64	0.32
GroupDRO *+ [38]	5.73	7.97	0.78	7.69	10.83	0.78	3.41	5.21	0.83	8.35	11.67	0.54	9.94	14.29	0.36
Coral *+ [53]	5.89	8.04	0.76	7.59	10.87	0.72	3.64	5.74	0.80	8.68	11.91	0.53	10.32	14.42	0.32
VREx *+ [18]	5.59	7.68	0.81	7.24	10.14	0.78	3.27	5.01	0.86	8.37	11.62	0.54	9.82	14.16	0.37
NCDG *+ [57]	5.31	7.56	0.82	7.32	10.35	0.77	3.12	5.16	0.85	8.47	11.81	0.52	10.14	14.46	0.34
Baseline * [36]	5.79	7.91	0.78	7.39	10.49	0.77	3.38	5.17	0.84	8.97	12.16	0.49	10.16	14.57	0.34
NEST *+ w/o \mathcal{L}_{CM}	4.92	7.19	0.83	6.91	9.94	0.78	2.79	4.62	0.89	8.21	11.56	0.55	9.58	13.58	0.38
NEST *+ w/o \mathcal{L}_{TA}	4.97	7.28	0.83	6.81	9.77	0.81	3.03	4.94	0.86	8.19	11.47	0.55	9.52	13.52	0.38
NEST *+ w/o \mathcal{L}_{DM}	4.89	7.11	0.84	6.87	9.81	0.79	2.93	4.77	0.88	8.14	11.43	0.56	9.49	13.49	0.40
NEST *+	4.67	6.79	0.86	6.71	9.59	0.81	2.88	4.69	0.89	7.86	11.15	0.58	9.27	13.79	0.41

Table 1. HR estimation results on MSDG protocol. It is worth noting that we have removed the GRU mechanism in Rhythmnet because it is an unfair comparison with using GRU to aggregate the temporal information of multiple clips instead of only one clip; * means that these methods use the STMap as the input of CNN and use the STMap augmentation strategy as mentioned in section 4.2; + means that these methods are based on baseline (Rhythmnet without GRU).

augmentation are used for STMap [5]. Then the STMap $x \in R^{25 \times 256 \times 3}$ is resized to $x' \in R^{64 \times 256 \times 3}$ as the input of the network. We use the ResNet-18 with removing the max pooling layer as the feature extractor, and the final full connection layer is used as the HR estimation head. At the same time, we designed a BVP signal regression head containing four blocks, each block involves a Conv., transpose, and two Conv. layers with batch normalization and Relu function. Hyper-parameter σ , ρ and τ are set to 5, 0.1, and 0.2 according to the experimental result, and the trade-off parameter $k_1 - k_5$ are set to 1, 0.1, 0.001, 0.1 and 0.01 based on the scale of losses. Adam optimizer with a learning rate of 0.001 is used for training. The batch size and iterations are set to 1024 and 20000.

Performance Metrics. Following existing methods [6, 28, 37, 50, 70], standard deviation (SD), mean absolute error (MAE), root mean square error (RMSE), and Pearson’s correlation coefficient (r) are used to evaluate the HR estimation. MAE, RMSE, and r are used to evaluate the HRV measurement, including low frequency (LF), high frequency (HF), and LF/HF.

4.3. Multi-Source Domain Generalization

4.3.1 HR Estimation

For remote HR estimation, all five datasets (i.e., UBFC [1], PURE [52], BUAA [68], VIPL [36], and V4V [47]) are used to evaluate the model generalization. Specifically, four datasets are used for training, and then the other dataset is

used to test. Firstly, we performance the iPhys² to evaluate three traditional method (i.e., GREEN [59], CHROM [9], and POS [63]) as shown in Tab. 1. Then, we reproduce five popular DL-based methods (i.e., DeepPhys [6], TS-CAN [25], Rhythmnet [36], and BVPNet [7], and Dual-GAN [28]) for comparison on the same setting. Subsequently, we directly applied the four DG methods (i.e., AD [12], GroupDRO [38], Coral [53], VREx [18], and NCDG [57]) to Rhythmnet with using DeepDG³.

The performance of traditional methods. As shown in Tab. 1, traditional algorithms show relatively acceptable results on some simple scenes, such as BUAA and UBFC, but very poor results on other datasets. It is difficult for these traditional algorithms to extract the physiological information on the V4V dataset, which shows that the traditional methods cannot deal with the noise caused by the complex environment and non-rigid motions.

The performance of DL-based methods. For DL-based methods, DeepPhys and TS-CAN even perform worse than traditional methods. Rhythmnet, BVPNet, and Dual-GAN perform much better than both other DL-based methods and traditional methods, which suggests that STMap is the more robust input of CNN than the aligned video. Compared with Rhythmnet, BVPNet and Dual-GAN have not significant improvement on the MSDG protocol. Because BVP-

²<https://github.com/danmcduff/iphs-toolbox>

³<https://github.com/jindongwang/transferlearning/tree/master/code/DeepDG>

Target	Method	LF-(u.n)			HF-(u.n)			LF/HF			HR-(bpm)		
		MAE↓	RMSE↓	r↑	MAE↓	RMSE↓	r↑	MAE↓	RMSE↓	r↑	MAE↓	RMSE↓	r↑
UBFC	GREEN [59]	0.2355	0.2841	0.0924	0.2355	0.2841	0.0924	0.6695	0.9512	0.0467	8.0184	9.1776	0.3634
	CHROM [9]	0.2221	0.2817	0.0698	0.2221	0.2817	0.0698	0.6708	1.0542	0.1054	7.2291	8.9224	0.5123
	POS [63]	0.2364	0.2861	0.1359	0.2364	0.2861	0.1359	0.6515	0.9535	0.1345	7.3539	8.0402	0.4923
	Baseline	0.0621	0.0813	0.1873	0.0621	0.0813	0.1873	0.1985	0.2667	0.3043	5.1542	7.4672	0.8165
	NEST	0.0597	0.0782	0.2017	0.0597	0.0782	0.2017	0.2138	0.2824	0.3179	4.7471	6.8876	0.8546
PURE	GREEN [59]	0.2539	0.3002	0.0326	0.2539	0.3002	0.0326	0.6525	0.8932	0.0417	10.3247	14.2693	0.4952
	CHROM [9]	0.2096	0.2751	0.1059	0.2096	0.2751	0.0759	0.5404	0.8266	0.1173	9.7914	12.7568	0.3732
	POS [63]	0.1959	0.2571	0.1684	0.1959	0.2571	0.1684	0.5373	0.846	0.1433	9.8273	13.4414	0.3432
	Baseline	0.0671	0.0923	0.6046	0.0671	0.0923	0.6046	0.2864	0.4184	0.5526	8.2542	11.1765	0.6832
	NEST	0.0635	0.0874	0.6422	0.0635	0.0874	0.6422	0.2255	0.3505	0.5734	7.6889	10.4783	0.7255
BUAA	GREEN [59]	0.3472	0.3951	0.0871	0.3472	0.3951	0.0871	0.6453	0.8632	0.0921	5.8231	7.9882	0.5624
	CHROM [9]	0.3786	0.3237	0.0682	0.3786	0.3237	0.0682	0.6813	0.8836	0.0715	6.0934	8.2938	0.5165
	POS [63]	0.3198	0.3762	0.0962	0.3198	0.3762	0.0962	0.6275	0.8424	0.1127	5.0407	7.1198	0.6374
	Baseline	0.1451	0.1681	0.2891	0.1451	0.1681	0.2891	0.5564	0.6904	0.2914	3.7852	6.3237	0.7492
	NEST	0.1436	0.1665	0.2955	0.1436	0.1665	0.2955	0.5514	0.6884	0.3004	3.3723	5.8806	0.7647

Table 2. HRV and HR estimation results on the MSDG protocol.

Net and Dual-GAN consider the performance of the single domain instead of focusing on solving domain shift. Therefore, Rhythmnet is also selected as a baseline to further analyze the generalization on MSDG protocol.

The performance of DG methods. To further improve the generalization of the model, we apply the existing popular DG methods to baseline (Rhythmnet) as shown in Tab. 1. There is a slight model degradation of AD [12], Coral [53], and GroupDRO [38] on some target datasets. The reason behind this is that these algorithms employ inaccurate domain labeling, specifically, a dataset often with different acquisition devices, lighting, and scene cannot be considered as the same domain. On the contrary, VREx [18] and NCDG [57] can slightly improve the performance of Rhythmnet without utilizing the domain label. Inaccurate domain division will make it difficult for the DG algorithm to effectively align the feature distribution across different domains, which is termed as *agnostic domain generalization*.

Without the domain labels, NEST experimentally manifested significant improvement on MSDG protocol. For the target dataset BUAA, the results without NEST-CM \mathcal{L}_{CM} show the best performance. Because NEST-CM may cause the model degradation with activating meaningless features at the instance level, especially for the BUAA with the slight domain shift.

4.3.2 HRV Estimation

We also conduct experiments for HRV (i.e., LF, HF, and LF/HF) and HR measurement on the PURE, UBFC, and BUAA. The reason for choosing only these three datasets is that VIPL and V4V do not have reliable ground-truth

Method	PURE			BUAA		
	MAE↓	RMSE↓	r↑	MAE↓	RMSE↓	r↑
GREEN [59]	10.32	14.27	0.52	5.82	7.99	0.56
CHROM [9]	9.79	12.76	0.37	6.09	8.29	0.51
POS [63]	9.82	13.44	0.34	5.04	7.12	0.63
Baseline	7.54	10.13	0.73	2.97	3.28	0.76
NEST	6.07	9.06	0.76	2.56	2.73	0.78

Table 3. HR estimation results on SSDG protocol by training on the UBFC-rPPG.

BVP signals, as mentioned in subsection 4.1. We also perform three tradition method (GREEN [59], CHMROM [9], POS [63]) by iPhys for comparison with our method as shown in Tab. 2. With a more robust NEST presentation distribution, our method also can get better performance than baseline and other traditional methods in all metrics.

4.4. Single-Source Domain Generalization

For a worst-case scenario, the model is trained on only one dataset to generalize to the unseen domains, namely single-source domain generalization (SSDG). Based on this, we select UBFC-rPPG as the source domain and then select PURE and BUAA as the target domain to further evaluate the generalization of NEST as shown in Tab. 3. As we can see, NEST achieves desired performance by training on only one dataset. The reason is that: (1) NEST-CM can avoid the appearance of abnormal NEST; (2) NEST-TA can find the shared and robust feature of the samples within the same label; (3) NEST-DM can distinguish the boundaries between different samples to improve the generalization.

Method	SD↓	MAE↓	RMSE↓	r↑
SAMC [58]	18.0	15.9	21.0	0.11
POS [63]	15.3	11.5	17.2	0.30
CHROM [9]	15.1	11.4	16.9	0.28
I3D [3]	15.9	12.0	15.9	0.07
DeepPhy [6]	13.6	11.0	13.8	0.11
BVPNet [8]	7.75	5.34	7.85	0.70
RhythmNet [36]	8.11	5.30	8.14	0.76
CVD [37]	7.92	5.02	7.97	0.79
Physformer [70]	7.74	4.97	7.79	0.78
Dual-GAN [28]	7.63	4.93	7.68	0.81
Baseline	8.04	5.21	8.07	0.77
NEST	7.49	4.76	7.51	0.84

Table 4. HR estimation results by our method and several state-of-the-art methods on the VIPL-HR database.




Figure 5. Impacts of the hyperparameter (a) ρ , (b) σ , and (c) τ of our proposed method.

4.5. Intra-Dataset Testing on VIPL

Our proposed method can also improve the effectiveness of the model on intra-dataset testing by reducing the harmful effects of abnormal disturbances. Following [28,36,37,70], we test NEST by using a 5-fold cross-validation protocol on VIPL-HR. Our method compared with ten baseline methods, including the traditional methods (SAMC [58], POS [63] and CHROM [9]) and the DL-based methods (I3D [3], DeepPhy [6], RhythmNet [36], CVD [37], BVPNet [7], Physformer [70], and Dual-GAN [28]). The results of these methods are directly from [28,70] as shown in Tab. 4.

As shown in Tab. 4, the proposed methods outperform all the methods under all metrics. Compared with the baseline, it also gets a noticeable promotion, i.e., 0.4 in MAE and 0.53 in RMSE. It is worth noting that the VIPL dataset has multiple complex scenes and recording devices and cannot be considered as a single domain, termed the agnostic domain problem. This proves that NEST can effectively alleviate the domain shift without domain labels.

4.6. Further Analysis

We conduct further experiments on MSDG protocol (testing on UBFC with training on the other four datasets).

Impact of ρ in NEST-CM. The hyperparameter ρ determines how many bases need to be strengthened by NEST-CM \mathcal{L}_{CM} . As illustrated in Fig. 5 (a), NEST could achieve smaller RMSE when $\rho = 0.1$, indicating that the bases with small eigenvalue require more attention and fine-tuning.




Figure 6. The pearson’s correlation coefficient of the NEST representation of the samples with different labels for the anchor HR 60.

Impact of σ in NEST-TA. The hyperparameter σ tradeoffs the contribution of each selected sample to form the synthesized prototype. As illustrated in Fig. 5, NEST could achieve the best performance with $\sigma = 5$, which shows that the appropriate fusion of the nearby sample is conducive to improving model generalization. The specific reason behind this is that: (1) the synthesized prototype based on multi-domain samples can be used to enhance the domain invariant feature by NEST-TA; (2) NEST-TA can effectively aggregate the label relationship into NEST presentation distribution, which is performed by improving the similarity of NEST representation between the adjacent label samples.

To further clarify the second reason, we calculate the Pearson’s correlation coefficient of the NEST representation of the samples with different labels for the anchor HR 60, as shown in Fig. 6. It is obvious that the distribution changes with NEST-TA loss are more smooth across the continuous HR labels, which allows OOD samples with slight domain shifts to predict nearby heart rates around the ground truth HR value instead of other abnormal HR values.

Impact of τ in NEST-DM. The hyperparameter τ determines the pushing strength of the different samples. With the smallest τ , NEST-DM \mathcal{L}_{DM} has the overemphasized constraints to distinguish the NEST representation between different samples, which causes a harmful impact on the discriminant ability of the model. Overall, NEST-DM \mathcal{L}_{DM} is an effective way to improve the generalization of model.

5. Conclusion

This paper presents the NEural STructure (NEST) modeling, a novel framework that addresses the domain generalization problem in rPPG measurement for the first time. Our framework contains three regularization terms defined in the view of neural structure, which improves the generalization capacity in the different aspects. All three terms do not rely on domain labels, which are especially suitable in rPPG, as defining domains is difficult in this task. To compare the generalization of the existing algorithm and our method, we establish the first DG protocol for this task. We experimentally manifested the substantial improvement of our method on both MSDG and SSDG protocols.

References

- [1] Serge Bobbia, Richard Macwan, Yannick Benezeth, Alamin Mansouri, and Julien Dubois. Unsupervised skin tissue segmentation for remote photoplethysmography. *Pattern Recognition Letters*, 124:82–90, 2019.
- [2] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In *Proc. ICCV*, 2017.
- [3] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proc. IEEE CVPR*, pages 6299–6308, 2017.
- [4] Beidi Chen, Tri Dao, Eric Winsor, Zhao Song, Atri Rudra, and Christopher Ré. Scatterbrain: Unifying sparse and low-rank attention. *NeurIPS*, 34:17413–17426, 2021.
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607. PMLR, 2020.
- [6] Weixuan Chen and Daniel McDuff. Deepphys: Video-based physiological measurement using convolutional attention networks. In *Proc. ECCV*, pages 349–365, 2018.
- [7] Abhijit Das, Hao Lu, Hu Han, Antitza Dantcheva, Shiguang Shan, and Xilin Chen. Bvpnet: Video-to-bvp signal prediction for remote heart rate estimation. In *FG*, pages 01–08. IEEE, 2021.
- [8] Abhijit Das, Hao Lu, Hu Han, Antitza Dantcheva, Shiguang Shan, and Xilin Chen. Bvpnet: Video-to-bvp signal prediction for remote heart rate estimation. In *FG*, pages 01–08. IEEE, 2021.
- [9] Gerard De Haan and Vincent Jeanne. Robust pulse rate from chrominance-based rppg. *IEEE Trans. Biomed. Eng.*, 60(10):2878–2886, 2013.
- [10] Gerard De Haan and Arno Van Leest. Improved motion robustness of remote-ppg by using the blood volume pulse signature. *Physiol. Meas.*, 35(9):1913, 2014.
- [11] Zhijie Deng, Yucen Luo, and Jun Zhu. Cluster alignment with a teacher for unsupervised domain adaptation. In *Proc. CVPR*, pages 9944–9953, 2019.
- [12] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189. PMLR, 2015.
- [13] Ming-Shiang Chen Gee-Sern Hsu, ArulMurugan Ambikapathi. Deep learning with time-frequency representation for pulse estimation. In *Proc. IJCB*, pages 642–650, 2017.
- [14] Min Hu, Fei Qian, Dong Guo, Xiaohua Wang, Lei He, and Fuji Ren. Eta-rppgnet: effective time-domain attention network for remote heart rate measurement. *IEEE Trans. on Instrum. and Meas.*, 70:1–12, 2021.
- [15] Yunpei Jia, Jie Zhang, Shiguang Shan, and Xilin Chen. Single-side domain generalization for face anti-spoofing. In *Proc. CVPR*, pages 8484–8493, 2020.
- [16] Xiang Jiang, Qicheng Lao, Stan Matwin, and Mohammad Havaei. Implicit class-conditioned domain alignment for unsupervised domain adaptation. In *ICML*, pages 4816–4827. PMLR, 2020.
- [17] Viktor Kessler, Patrick Thiam, Mohammadreza Amirian, and Friedhelm Schwenker. Pain recognition with camera photoplethysmography. In *IEEE IPTA*, pages 1–5, 2017.
- [18] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). In *ICML*, pages 5815–5826. PMLR, 2021.
- [19] Eugene Lee, Evan Chen, and Chen-Yi Lee. Meta-rppg: Remote heart rate estimation using a transductive meta-learner. In *ECCV*, pages 392–409. Springer, 2020.
- [20] Alexander Lehner, Stefano Gasperini, Alvaro Marcos-Ramiro, Michael Schmidt, Mohammad-Ali Nikouei Mahani, Nassir Navab, Benjamin Busam, and Federico Tombari. 3d-vfield: Adversarial augmentation of point clouds for domain generalization in 3d object detection. In *Proc. CVPR*, pages 17295–17304, 2022.
- [21] Magdalena Lewandowska, Jacek Rumiński, Tomasz Kocejko, and Jędrzej Nowak. Measuring pulse rate with a webcam—a non-contact method for evaluating cardiac activity. In *FedCSIS*, pages 405–410. IEEE, 2011.
- [22] Daiqing Li, Junlin Yang, Karsten Kreis, Antonio Torralba, and Sanja Fidler. Semantic segmentation with generative models: Semi-supervised learning and strong out-of-domain generalization. In *Proc. CVPR*, pages 8300–8311, 2021.
- [23] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Proc. AAAI*, 2018.
- [24] Chuang Lin, Zehuan Yuan, Sicheng Zhao, Peize Sun, Changhu Wang, and Jianfei Cai. Domain-invariant disentangled network for generalizable object detection. In *Proc. CVPR*, pages 8771–8780, 2021.
- [25] Xin Liu, Josh Fromm, Shwetak Patel, and Daniel McDuff. Multi-task temporal shift attention networks for on-device contactless vitals measurement. *NeurIPS*, 33:19400–19411, 2020.
- [26] Xin Liu, Ziheng Jiang, Josh Fromm, Xuhai Xu, Shwetak Patel, and Daniel McDuff. Metaphys: few-shot adaptation for non-contact physiological measurement. In *Proc. HIL*, pages 154–163, 2021.
- [27] Zhenhua Liu, Jizheng Xu, Xiulian Peng, and Ruiqin Xiong. Frequency-domain dynamic pruning for convolutional neural networks. *NeurIPS*, 31, 2018.
- [28] Hao Lu, Hu Han, and S Kevin Zhou. Dual-gan: Joint bvp and noise modeling for remote physiological measurement. In *Proc. IEEE CVPR*, pages 12404–12413, 2021.
- [29] Yawei Luo, Ping Liu, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Category-level adversarial adaptation for semantic segmentation using purified features. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2021.
- [30] Fangrui Lv, Jian Liang, Shuang Li, Bin Zang, Chi Harold Liu, Ziteng Wang, and Di Liu. Causality inspired representation learning for domain generalization. In *Proc. CVPR*, pages 8046–8056, 2022.
- [31] Lei Ma, Felix Juefei-Xu, Fuyuan Zhang, Jiyuan Sun, Minhui Xue, Bo Li, Chunyang Chen, Ting Su, Li Li, Yang Liu, et al. Deepgauge: Multi-granularity testing criteria for deep learning systems. In *Proc. ASE*, pages 120–131, 2018.

- [32] Daniel McDuff. Camera measurement of physiological vital signs. *CSUR*, 2021.
- [33] Daniel McDuff. Applications of camera-based physiological measurement beyond healthcare. In *CVSM*, pages 165–177. Elsevier, 2022.
- [34] Daniel McDuff, Sarah Gontarek, and Rosalind W Picard. Improvements in remote cardiopulmonary measurement using a five band digital camera. *IEEE Trans. Biomed. Eng.*, 61(10):2593–2601, 2014.
- [35] Xuesong Niu, Hu Han, Shiguang Shan, and Xilin Chen. Synrhythm: Learning a deep heart rate estimator from general to specific. In *Proc. IEEE ICPR*, pages 3580–3585, 2018.
- [36] Xuesong Niu, Shiguang Shan, Hu Han, and Xilin Chen. Rhythmnet: End-to-end heart rate estimation from face via spatial-temporal representation. *IEEE Trans. on Image Process.*, 29:2409–2423, 2020.
- [37] Xuesong Niu, Zitong Yu, Hu Han, Xiaobai Li, Shiguang Shan, and Guoying Zhao. Video-based remote physiological measurement via cross-verified feature disentangling. In *Proc. ECCV*, 2020.
- [38] Giambattista Parascandolo, Alexander Neitz, Antonio Orvieto, Luigi Gresle, and Bernhard Schölkopf. Learning explanations that are hard to vary. *ICLR2022*, 2022.
- [39] Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. Deepxplore: Automated whitebox testing of deep learning systems. In *Proc. SOSP*, pages 1–18, 2017.
- [40] Duo Peng, Yinjie Lei, Munawar Hayat, Yulan Guo, and Wen Li. Semantic-aware domain generalized segmentation. In *Proc. CVPR*, pages 2594–2605, 2022.
- [41] Vihari Piratla, Praneeth Netrapalli, and Sunita Sarawagi. Efficient domain generalization via common-specific low-rank decomposition. In *ICML*, pages 7728–7738. PMLR, 2020.
- [42] Ming-Zher Poh, Daniel J McDuff, and Rosalind W Picard. Non-contact, automated cardiac pulse measurements using video imaging and blind source separation. *Opt. Express*, 18(10):10762–10774, 2010.
- [43] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proc. CVPR*, pages 12556–12565, 2020.
- [44] Ying Qiu, Yang Liu, Juan Arteaga-Falconi, Haiwei Dong, and Abdulmotaleb El Saddik. Evm-cnn: Real-time contactless heart rate estimation from facial video. *IEEE Trans. Multimedia*, 2019.
- [45] Attila Reiss, Ina Indlekofer, Philip Schmidt, and Kristof Van Laerhoven. Deep ppg: large-scale heart rate estimation with convolutional neural networks. *Sensors*, 19(14):3079, 2019.
- [46] Jiahuan Ren, Zhao Zhang, Richang Hong, Mingliang Xu, Haijun Zhang, Mingbo Zhao, and Meng Wang. Robust low-rank convolution network for image denoising. In *Proc. ACM MM*, pages 6211–6219, 2022.
- [47] Ambareesh Revanur, Zhihua Li, Umur A Ciftci, Lijun Yin, and László A Jeni. The first vision for vitals (v4v) challenge for non-contact video-based physiological estimation. In *Proc. CVPR workshop*, pages 2760–2767, 2021.
- [48] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Sidhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. *ICLR*, 2018.
- [49] Rencheng Song, Huan Chen, Juan Cheng, Chang Li, Yu Liu, and Xun Chen. PulseGAN: Learning to generate realistic pulse waveforms in remote photoplethysmography. *IEEE J-BHI*, pages 1–1, 2021.
- [50] Rencheng Song, Senle Zhang, Chang Li, Yunfei Zhang, Juan Cheng, and Xun Chen. Heart rate estimation from facial videos using a spatiotemporal representation with convolutional neural networks. *IEEE Trans. Instrum. Meas.*, 2020.
- [51] Radim Špetlík, Vojtech Franc, and Jirí Matas. Visual heart rate estimation with convolutional neural network. In *Proc. BMVC*, pages 3–6, 2018.
- [52] Ronny Stricker, Steffen Müller, and Horst-Michael Gross. Non-contact video-based pulse rate measurement on a mobile service robot. In *Proc. IEEE ISRHIC*, pages 1056–1062, 2014.
- [53] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*, pages 443–450. Springer, 2016.
- [54] Bo Sun, Qinglan Wei, Liandong Li, Qihua Xu, Jun He, and Lejun Yu. Lstm for dynamic emotion and group emotion recognition in the wild. In *Proc. ACM ICMI*, pages 451–457, 2016.
- [55] Zhaodong Sun and Xiaobai Li. Contrast-phys: Unsupervised video-based remote physiological measurement via spatiotemporal contrast. In *Proc. ECCV*, 2022.
- [56] Ling Tang, Wen Shen, Zhanpeng Zhou, Yuefeng Chen, and Quanshi Zhang. Defects of convolutional decoder networks in frequency representation. *arXiv preprint arXiv:2210.09020*, 2022.
- [57] Chris Xing Tian, Haoliang Li, Xiaofei Xie, Yang Liu, and Shiqi Wang. Neuron coverage-guided domain generalization. *IEEE Trans. on Pattern Anal. Mach. Intell.*, 2022.
- [58] Sergey Tulyakov, Xavier Alameda-Pineda, Elisa Ricci, Lijun Yin, Jeffrey F Cohn, and Nicu Sebe. Self-adaptive matrix completion for heart rate estimation from face videos under realistic conditions. In *Proc. IEEE CVPR*, pages 2396–2404, 2016.
- [59] Wim Verkrusse, Lars O Svaasand, and J Stuart Nelson. Remote plethysmographic imaging using ambient light. *Opt. Express*, 16(26):21434–21445, 2008.
- [60] Tan Wad, Qianru Sun, Sugiri Pranata, Karlekar Jayashree, and Hanwang Zhang. Equivariance and invariance inductive bias for learning from insufficient data. In *ECCV*, pages 241–258. Springer, 2022.
- [61] Hao Wang, Euijoon Ahn, and Jinman Kim. Self-supervised representation learning framework for remote physiological measurement using spatiotemporal augmentation loss. In *ACM MM*, 2022.
- [62] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Tao Qin, Wang Lu, Yiqiang Chen, Wenjun Zeng, and Philip Yu. Generalizing to unseen domains: A survey on domain generalization. *IEEE Trans. on Knowl. Data Eng.*, 2022.
- [63] Wenjin Wang, Albertus C den Brinker, Sander Stuijk, and Gerard de Haan. Algorithmic principles of remote ppg. *IEEE Trans. Biomed. Eng.*, 64(7):1479–1491, 2017.

- [64] Wenjin Wang, Sander Stuijk, and Gerard De Haan. Exploiting spatial redundancy of image sensor for motion robust rppg. *IEEE Trans. Biomed. Eng.*, 62(2):415–425, 2015.
- [65] Yunhe Wang, Chang Xu, Shan You, Dacheng Tao, and Chao Xu. Cnnpack: Packing convolutional neural networks in the frequency domain. In *NeurIPS*, volume 29, 2016.
- [66] Zijian Wang, Yadan Luo, Ruihong Qiu, Zi Huang, and Mahsa Baktashmotlagh. Learning to diversify for single domain generalization. In *Proc. CVPR*, pages 834–843, 2021.
- [67] Zhi-Kuan Wang, Ying Kao, and Chiou-Ting Hsu. Vision-based heart rate estimation via a two-stream cnn. In *Proc. IEEE ICIP*, pages 3327–3331, 2019.
- [68] Lin Xi, Weihai Chen, Changchen Zhao, Xingming Wu, and Jianhua Wang. Image enhancement for remote photoplethysmography in a low-light environment. In *FG*, pages 1–7. IEEE, 2020.
- [69] Zitong Yu, Wei Peng, Xiaobai Li, Xiaopeng Hong, and Guoying Zhao. Remote heart rate measurement from highly compressed facial videos: an end-to-end deep learning solution with video enhancement. In *Proc. IEEE ICCV*, pages 151–160, 2019.
- [70] Zitong Yu, Yuming Shen, Jingang Shi, Hengshuang Zhao, Philip Torr, and Guoying Zhao. Physformer: Facial video-based physiological measurement with temporal difference transformer. *IEEE CVPR*, 2022.
- [71] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *Proc. CVPR*, pages 2100–2110, 2019.