

TransRPPG: Remote Photoplethysmography Transformer for 3D Mask Face Presentation Attack Detection

Zitong Yu, Xiaobai Li, Pichao Wang and Guoying Zhao, *Senior Member, IEEE*

Abstract—3D mask face presentation attack detection (PAD) plays a vital role in securing face recognition systems from emergent 3D mask attacks. Recently, remote photoplethysmography (rPPG) has been developed as an intrinsic liveness clue for 3D mask PAD without relying on the mask appearance. However, the rPPG features for 3D mask PAD are still needed expert knowledge to design manually, which limits its further progress in the deep learning and big data era. In this letter, we propose a pure rPPG transformer (TransRPPG) framework for learning intrinsic liveness representation efficiently. At first, rPPG-based multi-scale spatial-temporal maps (MSTmap) are constructed from facial skin and background regions. Then the transformer fully mines the global relationship within MSTmaps for liveness representation, and gives a binary prediction for 3D mask detection. Comprehensive experiments are conducted on two benchmark datasets to demonstrate the efficacy of the TransRPPG on both intra- and cross-dataset testings. Our TransRPPG is lightweight and efficient (with only 547K parameters and 763M FLOPs), which is promising for mobile-level applications.

Index Terms—rPPG, transformer, 3D mask, face presentation attack detection, spatio-temporal map.

I. INTRODUCTION

FACE recognition [1] technology has been widely used in many interactive intelligent systems due to their convenience and remarkable accuracy. However, face recognition systems are still vulnerable to presentation attacks (PAs) ranging from print [2], replay [3] and emerging 3D mask [4] attacks. Therefore, both the academia and industry have realized the critical role of face presentation attack detection (PAD) [5] technique for securing the face recognition system.

In the past decade, both traditional [6]–[8] and deep learning-based [9]–[14] methods have shown effectiveness for face PAD especially for 2D (e.g., print and replay) attacks. On one hand, most previous approaches extract the static appearance based clues (e.g., color texture [15] and noise artifacts [12]) to distinguish the PAs from the bonafide. On the other hand, a few methods exploit the dynamic inconsistency (e.g., motion blur [3] and temporal depth [16]) between the

bonafide and PAs for discrimination. Benefited from the above-mentioned advanced techniques, 2D PA detection has made satisfied progress under most normal scenarios.

Meanwhile, 3D mask attack has attracted increasing attention since the customized 3D mask can be easily made at an affordable price [17]. Despite strong ability on detecting low-quality 3D print mask, appearance-based methods still suffer from performance drops when encountering high-fidelity 3D masks with fine-grained texture and shape as real faces [18]. In contrast, some researchers devote to exploring the feasibility to utilize appearance-independent liveness clues (e.g., remote photoplethysmography (rPPG) [19]–[22]) for 3D mask PAD instead of the traditional texture/motion-based patterns.

rPPG [23] is a new technique to recover the physiological signals under ambient light [24] via analyzing the skin color changes caused by periodic cardiac heartbeats. In the early stage, most methods [24]–[28] capture subtle color changes on particular facial regions of interest (ROI). Then some end-to-end deep learning methods [29]–[33] are proposed to recover rPPG signals from facial videos directly. Recently, learning upon the constructed spatio-temporal map (STmap) [34], [35] consisting of raw color variance signals from multiple facial ROIs, are proven to learn rPPG features more efficiently.

One nature question is that why rPPG is suitable for 3D mask face PAD? On one hand, rPPG can be captured using only a simple RGB camera under ambient light, which satisfies most of the video recording conditions in 3D mask face PAD. On the other hand, live/periodic pulse signals can only be observed on genuine faces but not on masked faces because the 3D mask blocks the light transmission from the facial skin [19]. As the rPPG clues are independent to the mask appearance, the rPPG-based 3D mask PAD methods [19], [21], [22] can detect the high-fidelity mask well and shows good generalization capacity.

To alleviate the influence of environmental noise, existing 3D mask PAD approaches [19], [22] usually design complex hand-crafted rPPG features (e.g., cross-correlation spectrum). In other words, how to efficiently represent rPPG features for 3D mask face PAD in a data-driven fashion is still unexplored. Recently, transformer [22] shows its strong long-range relationship modeling ability, and achieves state-of-the-art performance in both natural language processing (NLP) [36] and computer vision (CV) [37]–[39] tasks. Inspired by the rich global attention characteristic of the vision transformer (ViT) [38], in this letter we propose a pure rPPG transformer framework, named TransRPPG, for learning intrinsic live-

Manuscript received April xx, 2021. This work was supported by the Academy of Finland for project MiGA (grant 316765), ICT 2023 project (grant 328115), Infotech Oulu, project 6+E (grant 323287) funded by Academy of Finland, and project PhInGAIN (grant 200414) funded by The Finnish Work Environmental Fund. The authors wish to acknowledge CSC-IT Center for Science, Finland, for computational resources. (Corresponding author: Guoying Zhao)

Z. Yu, X. Li and G. Zhao are with the Center for Machine Vision and Signal Analysis, University of Oulu, Oulu 90014, Finland. E-mail: {zitong.yu, xiaobai.li, guoying.zhao}@oulu.fi.

P. Wang is with Alibaba Group, Bellevue, WA, 98004, USA. E-mail: pichao.wang@alibaba-inc.com

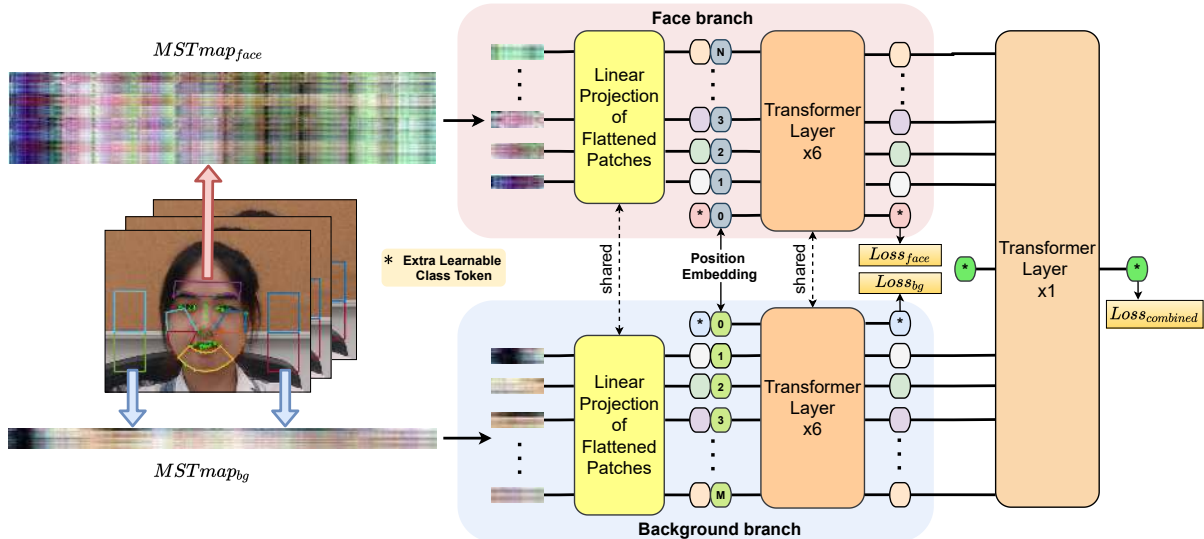


Fig. 1: Overview of TransRPPG for 3D mask face PAD. Given a face video, two MSTmaps are constructed from facial and background region, respectively. Then two-branch vision transformers (shared) are used to extract the respective rPPG and environmental features. Finally, the combined features are refined via an extra transformer for binary (bonafide or mask attack) prediction.

ness representation efficiently. Similar to the work [35], we first construct rPPG-based multi-scale spatial-temporal maps (MSTmap) on facial skin and background regions as the input of vision transformer [38]. Then the transformer fully mines the global relationship within MSTmap for liveness representation, and gives a binary prediction for 3D mask detection. Our contribution includes:

- We propose the first pure transformer framework (TranRPPG) for automatic rPPG feature representation on 3D mask face PAD task. We also explore various task-aware inputs, architectures, and loss functions for TranRPPG, which is insightful for both rPPG and 3D mask face PAD communities.
- We conduct intra- and cross-dataset tests and show that the TranRPPG achieves superior or on par state-of-the-art performance on 3D mask face PAD. Moreover, our TransRPPG is lightweight and efficient (with only 547K parameters and 763M FLOPs), which is promising for mobile-level applications.

II. METHODOLOGY

In this section, we first give details about the multi-scale spatial-temporal map generation. Then we introduce the transformer-based backbone for 3D mask face PAD. An overview of the proposed method is illustrated in Fig. 1.

A. Multi-scale Spatial Temporal Map Generation

To explicitly utilize sufficient rPPG signals from global and local regions for 3D mask face PAD, we follow the similar manner about multi-scale spatial temporal maps (MSTmap) generation in [35]. As shown in Fig. 1 (left), We first use an open source face detector OpenFace [40] to detect the facial landmarks, based on which we define the most informative region of interest (ROIs) for physiological measurement in face, i.e., the forehead and cheek areas. Specifically, for the t -th frame of a video, we first get a set of n informative

regions of face $R_t = \{R_{1t}, R_{2t}, \dots, R_{nt}\}$. Then we calculate the average RGB pixel values of each color channel for all the non-empty subsets of R_t , which are $2^n - 1$ combinations of the elements in R_t . For each video with T frames, the $2^n - 1$ temporal signals of each channel are placed into rows, and we can get the face MSTmap $MSTmap_{face}$ with the size of $(2^n - 1) \times T \times 3$ for each video. Besides raw signals from RGB channels, we also consider $MSTmap_{face}$ from YUV color space or other rPPG signal measurement approaches [28], [41], which will be discussed in Sec. III-C.

Besides the possible skin color changes in facial regions, the environmental patterns from background regions are also helpful for mask PAD. On one hand, the environmental light clues from the background benefit the recovery of the robust facial rPPG signals. On the other hand, the latent noise (e.g., camera noise) from the background can be treated as a contrast with facial signals (large difference for the bonafide while similarity for mask attacks) for liveness discrimination. Therefore, as shown in Fig. 1 (left), we also define m background regions on both sides of the face to construct the background MSTmap $MSTmap_{bg}$ with the size $(2^m - 1) \times T \times 3$ for each video.

B. rPPG Transformer for 3D Mask Face PAD

Image Sequentialization. Similar to ViT [38], we first preprocess the input face and background MSTmaps into a sequence of flattened patches X_{face} and X_{bg} , respectively. However, the original split method cuts the images into non-overlapping patches, which discards discriminative information of some local neighboring signal clips or informative interested regions. To alleviate this problem, we propose to generate overlapping patches with sliding window. Specifically, with the targeted patch size $P_H \times P_W$ and step size $S_H \times S_W$ of sliding window, the input face MSTmap with size $H \times W$ will be split into N patches where

$$N = N_H \times N_W = \left\lfloor \frac{H - P_H + S_H}{S_H} \right\rfloor \times \left\lfloor \frac{W - P_W + S_W}{S_W} \right\rfloor. \quad (1)$$

Similarly, the background MSTmap is partitioned into M patches. This simple yet fine-grained patch partition strategy improves the performance remarkably.

Patch Embedding. We map the vectorized patches X_{face} into a latent D -dimensional embedding space using a trainable linear projection. Then, the patch embeddings Z_{face}^0 can be formulated with an additional learnable position embedding to retain positional information:

$$Z_{face}^0 = [X_{faceCls}; X_{face}^1 \mathbf{E}, X_{face}^2 \mathbf{E}, \dots, X_{face}^N \mathbf{E}] + \mathbf{E}_{pos}, \quad (2)$$

where $\mathbf{E} \in \mathbb{R}^{(P^2 \cdot C) \times D}$ is the patch embedding projection from the original C -channel space, and $\mathbf{E}_{pos} \in \mathbb{R}^{(N+1) \times D}$ denotes the position embedding. The class token $X_{faceCls}$, i.e., a learnable embedding, is concatenated to the patch embeddings. Similarly, the background embedding Z_{bg}^0 can be extracted using the shared linear projection \mathbf{E} but independent class token X_{bgCls} and position embedding \mathbf{E}'_{pos} .

Transformer Encoder. The Transformer encoder contains L layers of multihead self-attention (MSA) and multi-layer perceptron (MLP) blocks. Thus the output of the l -th layer can be written as follows:

$$\begin{aligned} Z^l &= MSA(LN(Z^{l-1})) + Z^{l-1}, \quad l \in 1, 2, \dots, L \\ Z^l &= MLP(LN(Z^l)) + Z^l, \quad l \in 1, 2, \dots, L \end{aligned} \quad (3)$$

where $Z^l \in \mathbb{R}^{(N+1) \times D}$ is the token features in the l -th layer. $LN(\cdot)$ denotes the layer normalization operation, and MSA is composed of h parallel self-attention (SA),

$$\begin{aligned} [Q, K, V] &= Z^l U_{QKV}, \\ SA(Z^l) &= Softmax(QK^T / \sqrt{D'}) V, \end{aligned} \quad (4)$$

where $U_{QKV} \in \mathbb{R}^{D \times 3D'}$ is the weight matrix for linear transformation, and $A = Softmax(QK^T / \sqrt{D'})$ is the attention map. The output of MSA is the concatenation of h attention head outputs

$$MSA(Z^l) = [SA_1(Z^l); SA_2(Z^l); \dots; SA_h(Z^l)] U_{MSA}, \quad (5)$$

where $U_{MSA} \in \mathbb{R}^{hD' \times D}$. As illustrated in Fig. 1, the two transformer branches (i.e., face and background) first project the patch embeddings Z_{face}^0 and Z_{bg}^0 , and then L -layer shared transformer encoder is utilized for global attentional features (Z_{face}^L and Z_{bg}^L) representation. Subsequently, all token features from Z_{face}^L and Z_{bg}^L (except the respective class token features) as well as an joint class token X_{comCls} are concatenated, and an extra transformer layer is utilized for face and background context aggregation.

Hierarchical Supervision. In order to provide explicit supervision signals for TransRPPG, we design three Binary Cross Entropy (BCE) losses (bonafide vs. mask attack) for face, background, and combined branches, respectively. To be specific, three class tokens $X_{faceCls}$, X_{bgCls} , and X_{comCls} with independent linear classification heads are supervised with BCE losses \mathcal{L}_{face} , \mathcal{L}_{bg} , and $\mathcal{L}_{combined}$, respectively. The overall loss can be formulated as $\mathcal{L}_{overall} = \mathcal{L}_{face} + \mathcal{L}_{bg} + \mathcal{L}_{combined}$. As there are usually no liveness clues in the background regions, the groundtruth for \mathcal{L}_{bg} are all simply regarded as ‘mask attack’.




Fig. 2: Face samples from two 3D mask face PAD datasets: (a) 3DMAD [17], and (b) HKBU-MARsV2 [4].

III. EXPERIMENTS

A. Datasets and Metrics

Two public datasets (see Fig. 2) are employed in our experiments. The **3DMAD** [17] dataset contains 255 videos from 17 subjects with 17 low-fidelity hard resin masks, which are recorded at 640×480 , 30fps under controlled lighting condition. The **HKBU-MARsV2** [4] dataset consists of 1008 videos from 12 subjects with 6 low- and 6 high-fidelity hard resin masks, which are recorded at 1280×720 , 25fps under variant room lights. In terms of the intra-dataset testings on 3DMAD, MARsV2, and combined dataset (3DMAD+MARsV2), leave-one-subject-out cross-validation protocol is used. As for the cross-dataset testings between 3DMAD and MARsV2, all videos in one dataset are used for training while those in the other one for testing. Area Under the Curve (AUC), Equal Error Rate (EER), Half Total Error Rate (HTER) [17], and False Fake Rate (FFR) when False Liveness Rate (FLR) equals to 0.01 are used as evaluation metrics.

B. Implementation Details

All 10-second videos are linearly interpolated into 30 fps to keep the same frame numbers ($T = 300$). There are $n = 6$ facial regions and $m = 4$ background regions for two MSTmaps constructions, which results in $MSTmap_{face}$ and $MSTmap_{bg}$ with size $63 \times 300 \times 3$ and $15 \times 300 \times 3$, respectively. In terms of transformer settings, patch size $P_H = 3$, $P_W = 30$, step size $S_H = 1$, $S_W = 15$, (hidden) embedding channels $D = D' = 96$, heads number = 3, and $L = 6$ layers are utilized. All MLP blocks has two fully-connected layers, where feature dimension doubles in the first layer. Our proposed method is implemented with Pytorch. The models are trained with Adam optimizer and the initial learning rate (lr) and weight decay are $1e-4$ and $5e-5$, respectively. We train models on a single V100 GPU with batchsize 10 for maximum 60 epochs while lr halves in the 45th epoch.

C. Ablation Study

All ablation studies are intra-dataset tested on the 3DMAD dataset. For convenience, single face branch of TransRPPG is first adopted as the backbone in experiments, and we will study the effects of two-branch framework at last.

Impact of Color Space for MSTmap. Fig. 3(a) shows that MSTmaps construction from raw RGB channels are best-suited for TransRPPG. Despite great rPPG measurement performance in previous works (e.g., G channel [26], CHROM [27], POS [28], and RGB+YUV [35]), these color spaces cannot provide extra improvement for TransRPPG.

Impact of Patch Size. As shown in Fig. 3(b), patch size of STMmap tokenization influences liveness representation a

TABLE I: Intra-dataset results on 3DMAD, MARsV2, and combined (3DMAD+MARsV2) datasets. Best results are marked in **bold** and second best in underline. The rPPG-based methods are marked with \star while appearance-based methods with \blacktriangle .

Method	3DMAD			MARsV2			Combined		
	EER(%) \downarrow	AUC(%) \uparrow	FFR@FLR=0.01 \downarrow	EER(%) \downarrow	AUC(%) \uparrow	FFR@FLR=0.01 \downarrow	EER(%) \downarrow	AUC(%) \uparrow	FFR@FLR=0.01 \downarrow
MS-LBP [17] \blacktriangle	2.71	99.7	3.62	22.5	85.8	95.1	16.6	91.0	64.2
CTA [42] \blacktriangle	4.24	<u>99.3</u>	<u>11.8</u>	23.0	82.3	89.2	18.9	87.7	95.7
VGG16 [43] \blacktriangle	6.63	98.9	18.5	15.2	91.4	93.5	14.5	93.5	71.5
GrPPG [19] \star	14.4	92.2	36.0	16.4	89.4	32.9	15.2	91.1	42.8
LrPPG [4] \star	9.64	95.5	14.8	9.07	<u>97.0</u>	38.9	9.21	95.7	29.4
CFrPPG [22] \star	7.44	96.8	13.6	4.04	99.3	17.8	<u>6.54</u>	97.6	15.5
TransRPPG (Ours)\star	<u>2.38</u>	<u>98.77</u>	13.73	<u>8.47</u>	96.82	<u>29.79</u>	5.93	97.95	<u>22.46</u>

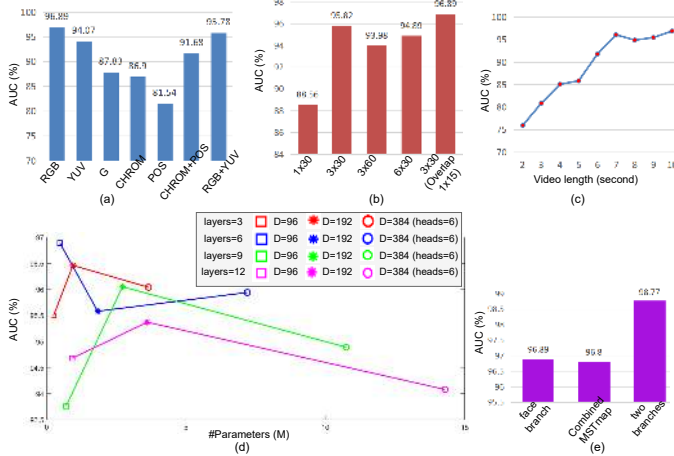


Fig. 3: Ablation studies on (a) Color space for MSTmap; (b) Patch size for image sequentialization; (c) Video length; (d) Transformers' depth/width; and (e) Background branch.

lot. When partition with size 1×30 ($P_H \times P_W$), i.e., only single spatial region clues within each patch, the AUC drops sharply. In contrast, large spatial (e.g., $P_H = 6$) or temporal (e.g., $P_W = 60$) patch sizes could not bring extra improvements. Patch size with 3×30 performs the best while introducing patch overlapping 1×15 ($S_H \times S_W$) sampling boosts 1% AUC.

Impact of Video Length. Fig. 3(c) illustrates that AUC ascends reasonably when video length increases because of richer temporal context for global modeling. It is interesting to find that using 7s video clips could still achieve comparable performance (96.06% vs 96.89% AUC) as 10s ones. We also find the limitations that TransRPPG under time-constrained scenarios (< 5 s) have poor performance ($< 90\%$ AUC).

Impact of Transformers' Depth/Width. It is necessary to investigate the impact of transformers' depth/width for mask PAD. As illustrated in Fig. 3(d), layer number (depth) plays more important role than embedding dimension (width). TransRPPG with shallow layers (e.g., layer = 3 and 6) performs better than those with deeper layers (e.g., layer = 9 and 12). The highest AUC is achieved when $L = 6$ with $D = 92$.

Effects of the Background Branch. Fig. 3(e) illustrates the evaluation results w/ and w/o $MSTmap_{bg}$. There are two configurations with background map: 1) using combined MSTmap via concatenated $MSTmap_{face}$ and $MSTmap_{bg}$ for single face branch input; and 2) two-branch framework in Fig. 1. It is surprised that combined MSTmap performs even worse than $MSTmap_{face}$, which might be caused by the inharmonious attention learning within compound maps. In contrast, disentangled learning from two branches improves the AUC by nearly 2%. We also find that the constraints \mathcal{L}_{face}

TABLE II: Cross-dataset results on 3DMAD and MARsV2.

Method	3DMAD \rightarrow MARsV2		MARsV2 \rightarrow 3DMAD	
	AUC(%) \uparrow	FFR@FLR=0.01 \downarrow	AUC(%) \uparrow	FFR@FLR=0.01 \downarrow
MS-LBP [17] \blacktriangle	60.4	100.0	75.3	87.8
CTA [42] \blacktriangle	62.1	98.3	60.5	96.5
VGG16 [43] \blacktriangle	54.6	97.9	58.6	99.3
GrPPG [19] \star	86.7	78.5	87.2	94.5
LrPPG [21] \star	95.6	61.7	92.3	48.7
CFrPPG [22] \star	99.0	19.6	98.0	12.4
TransRPPG (Ours)\star	91.3	47.6	98.3	18.5

and \mathcal{L}_{bg} are helpful for liveness feature learning, improving the AUC from 97.51% to 98.77%.

D. Intra-dataset Testing

In this subsection, we compare our TransRPPG with three previous rPPG-based mask PAD methods (GrPPG [19], LrPPG [21], and CFrPPG [22]) as well as appearance-based baselines on 3DMAD, MARsV2, and combined datasets. As shown in the upper block of Tab. I, appearance-based methods perform well on 3DMAD but poorly on MARsV2 because the 3D mask attacks on the latter one are with higher fidelity and harder to distinguish via texture clues. In contrast, four rPPG-based approaches shown in the lower block have consistent good performance on both 3DMAD and MARsV2 datasets, indicating that the learned rPPG-based liveness features are independent from mask fidelity. Instead of LrPPG and CFrPPG using complex spectrum features, our TransRPPG uses simple time-domain signals, and achieves the best EER and AUC on the combined dataset. Introducing frequency-domain representation to TransRPPG might be a possible future direction.

E. Cross-dataset Testing

Here we alternatively train and test between 3DMAD and MARsV2 to validate the models' generalization capacity. As shown in Tab. II, it is clear that the proposed TransRPPG generalizes well from MARsV2 to 3DMAD because the training samples in MARsV2 are more sufficient and diverse to alleviate the overfitting problem. On the contrary, TransRPPG suffers from obvious performance drops when trained on 3DMAD and tested on MARsV2 due to the small amount of training data under single scenario. LrPPG and CFrPPG generalize well in cross testings, implying the importance of spectrum representation lacked in TransRPPG.

IV. CONCLUSION

In this letter, we propose a lightweight remote photoplethysmography transformer (TransRPPG) for 3D mask face presentation attack (PA) detection based on the facial and background multiscale spatio-temporal maps. In the future, we plan to explore TransRPPG on 1) detecting other PA types such as print, replay, and makeup; and 2) various rPPG-based applications like heart rate and stress estimation.

REFERENCES

- [1] J. Guo, X. Zhu, C. Zhao, D. Cao, Z. Lei, and S. Z. Li, "Learning meta face recognition in unseen domains," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 6163–6172.
- [2] Z. Zhang, J. Yan, S. Liu, Z. Lei, D. Yi, and S. Z. Li, "A face anti-spoofing database with diverse attacks," in *2012 5th IAPR international conference on Biometrics (ICB)*. IEEE, 2012, pp. 26–31.
- [3] L. Li, Z. Xia, A. Hadid, X. Jiang, H. Zhang, and X. Feng, "Replayed video attack detection based on motion blur analysis," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 9, pp. 2246–2261, 2019.
- [4] S. Liu, P. C. Yuen, S. Zhang, and G. Zhao, "3d mask face anti-spoofing with remote photoplethysmography," in *European Conference on Computer Vision*. Springer, 2016, pp. 85–100.
- [5] A. Liu, X. Li, J. Wan, Y. Liang, S. Escalera, H. J. Escalante, M. Madadi, Y. Jin, Z. Wu, X. Yu *et al.*, "Cross-ethnicity face anti-spoofing recognition challenge: A review," *IET Biometrics*, vol. 10, no. 1, pp. 24–43, 2021.
- [6] D. Wen, H. Han, and A. K. Jain, "Face spoof detection with image distortion analysis," *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 4, pp. 746–761, 2015.
- [7] Z. Boulkenafet, J. Komulainen, and A. Hadid, "Face antispoofing using speeded-up robust features and fisher vector encoding," *IEEE Signal Processing Letters*, vol. 24, no. 2, pp. 141–145, 2016.
- [8] K. Patel, H. Han, and A. K. Jain, "Secure face unlock: Spoof detection on smartphones," *IEEE transactions on information forensics and security*, vol. 11, no. 10, pp. 2268–2283, 2016.
- [9] Z. Yu, C. Zhao, Z. Wang, Y. Qin, Z. Su, X. Li, F. Zhou, and G. Zhao, "Searching central difference convolutional networks for face anti-spoofing," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 5295–5305.
- [10] Z. Yu, X. Li, X. Niu, J. Shi, and G. Zhao, "Face anti-spoofing with human material perception," in *European Conference on Computer Vision*. Springer, 2020, pp. 557–575.
- [11] Y. Liu, A. Jourabloo, and X. Liu, "Learning deep models for face anti-spoofing: Binary or auxiliary supervision," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 389–398.
- [12] A. Jourabloo, Y. Liu, and X. Liu, "Face de-spoofing: Anti-spoofing via noise modeling," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 290–306.
- [13] Z. Yu, J. Wan, Y. Qin, X. Li, S. Z. Li, and G. Zhao, "Nas-fas: Static-dynamic central difference network search for face anti-spoofing," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pp. 1–1, 2020.
- [14] Z. Yu, X. Li, J. Shi, Z. Xia, and G. Zhao, "Revisiting pixel-wise supervision for face anti-spoofing," *IEEE Transactions on Biometrics, Behavior, and Identity Science (TBIOM)*, 2021.
- [15] Z. Boulkenafet, J. Komulainen, and A. Hadid, "Face spoofing detection using colour texture analysis," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 8, pp. 1818–1830, 2016.
- [16] Z. Wang, Z. Yu, C. Zhao, X. Zhu, Y. Qin, Q. Zhou, F. Zhou, and Z. Lei, "Deep spatial gradient and temporal depth learning for face anti-spoofing," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 5042–5051.
- [17] N. Erdogmus and S. Marcel, "Spoofing face recognition with 3d masks," *IEEE transactions on information forensics and security*, vol. 9, no. 7, pp. 1084–1097, 2014.
- [18] S. Jia, G. Guo, and Z. Xu, "A survey on 3d mask presentation attack detection and countermeasures," *Pattern Recognition*, vol. 98, p. 107032, 2020.
- [19] X. Li, J. Komulainen, G. Zhao, P.-C. Yuen, and M. Pietikainen, "Generalized face anti-spoofing by detecting pulse from face videos," in *2016 23rd International Conference on Pattern Recognition (ICPR)*. IEEE, 2016, pp. 4244–4249.
- [20] B. Lin, X. Li, Z. Yu, and G. Zhao, "Face liveness detection by rppg features and contextual patch-based cnn," in *Proceedings of the 2019 3rd International Conference on Biometric Engineering and Applications*, 2019, pp. 61–68.
- [21] S. Liu, P. C. Yuen, S. Zhang, and G. Zhao, "3d mask face anti-spoofing with remote photoplethysmography," in *European Conference on Computer Vision*. Springer, 2016, pp. 85–100.
- [22] S.-Q. Liu, X. Lan, and P. C. Yuen, "Remote photoplethysmography correspondence feature for 3d mask face presentation attack detection," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 558–573.
- [23] X. Chen, J. Cheng, R. Song, Y. Liu, R. Ward, and Z. J. Wang, "Video-based heart rate measurement: Recent advances and future prospects," *IEEE Transactions on Instrumentation and Measurement*, vol. 68, no. 10, pp. 3600–3615, 2018.
- [24] W. Verkruyse, L. O. Svaasand, and J. S. Nelson, "Remote plethysmographic imaging using ambient light," *Optics express*, vol. 16, no. 26, pp. 21 434–21 445, 2008.
- [25] M.-Z. Poh, D. J. McDuff, and R. W. Picard, "Advancements in non-contact, multiparameter physiological measurements using a webcam," *IEEE transactions on biomedical engineering*, vol. 58, no. 1, pp. 7–11, 2010.
- [26] X. Li, J. Chen, G. Zhao, and M. Pietikainen, "Remote heart rate measurement from face videos under realistic situations," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 4264–4271.
- [27] G. De Haan and V. Jeanne, "Robust pulse rate from chrominance-based rppg," *IEEE Transactions on Biomedical Engineering*, vol. 60, no. 10, pp. 2878–2886, 2013.
- [28] W. Wang, A. C. den Brinker, S. Stuijk, and G. De Haan, "Algorithmic principles of remote ppg," *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 7, pp. 1479–1491, 2016.
- [29] R. Špetlík, V. Franc, and J. Matas, "Visual heart rate estimation with convolutional neural network," in *Proceedings of the British Machine Vision Conference, Newcastle, UK*, 2018, pp. 3–6.
- [30] W. Chen and D. McDuff, "Deepphys: Video-based physiological measurement using convolutional attention networks," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 349–365.
- [31] Z. Yu, X. Li, and G. Zhao, "Remote photoplethysmograph signal measurement from facial videos using spatio-temporal networks," in *Proceedings of the British Machine Vision Conference, Cardiff, UK*, 2019, pp. 1–12.
- [32] Z. Yu, W. Peng, X. Li, X. Hong, and G. Zhao, "Remote heart rate measurement from highly compressed facial videos: an end-to-end deep learning solution with video enhancement," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 151–160.
- [33] Z. Yu, X. Li, X. Niu, J. Shi, and G. Zhao, "Autohr: A strong end-to-end baseline for remote heart rate measurement with neural searching," *IEEE Signal Processing Letters*, vol. 27, pp. 1245–1249, 2020.
- [34] X. Niu, S. Shan, H. Han, and X. Chen, "Rhythmnet: End-to-end heart rate estimation from face via spatial-temporal representation," *IEEE Transactions on Image Processing*, vol. 29, pp. 2409–2423, 2019.
- [35] X. Niu, Z. Yu, H. Han, X. Li, S. Shan, and G. Zhao, "Video-based remote physiological measurement via cross-verified feature disentangling," in *European Conference on Computer Vision*. Springer, 2020, pp. 295–310.
- [36] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [37] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *European Conference on Computer Vision*. Springer, 2020, pp. 213–229.
- [38] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations (ICLR)*, 2021.
- [39] S. He, H. Luo, P. Wang, F. Wang, H. Li, and W. Jiang, "Transreid: Transformer-based object re-identification," *arXiv preprint arXiv:2102.04378*, 2021.
- [40] T. Baltrusaitis, A. Zadeh, Y. C. Lim, and L.-P. Morency, "Openface 2.0: Facial behavior analysis toolkit," in *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*. IEEE, 2018, pp. 59–66.
- [41] W. Wang, S. Stuijk, and G. De Haan, "Exploiting spatial redundancy of image sensor for motion robust rppg," *IEEE transactions on Biomedical Engineering*, vol. 62, no. 2, pp. 415–425, 2014.
- [42] Z. Boulkenafet, J. Komulainen, and A. Hadid, "Face anti-spoofing based on color texture analysis," in *2015 IEEE international conference on image processing (ICIP)*. IEEE, 2015, pp. 2636–2640.
- [43] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.