

# Disentangled Representation with Dual-stage Feature Learning for Face Anti-spoofing

Yu-Chun Wang\*, Chien-Yi Wang<sup>†</sup>, Shang-Hong Lai\*<sup>†</sup>

\*National Tsing Hua University, <sup>†</sup>Microsoft AI R&D Center, Taiwan

yuchun@gapp.nthu.edu.tw chiwa@microsoft.com shlai@microsoft.com

## Abstract

As face recognition is widely used in diverse security-critical applications, the study of face anti-spoofing (FAS) has attracted more and more attention. Several FAS methods have achieved promising performance if the attack types in the testing data are included in the training data, while the performance significantly degrades for unseen attack types. It is essential to learn more generalized and discriminative features to prevent overfitting to pre-defined spoof attack types. This paper proposes a novel dual-stage disentangled representation learning method that can efficiently untangle spoof-related features from irrelevant ones. Unlike previous FAS disentanglement works with one-stage architecture, we found that the dual-stage training design can improve the training stability and effectively encode the features to detect unseen attack types. Our experiments show that the proposed method provides superior accuracy than the state-of-the-art methods on several cross-type FAS benchmarks.

## 1. Introduction

Face recognition technologies have been frequently applied in our daily life in recent years. The associated facial recognition security issues become critical because hackers may easily break into the face verification system by different presentation attacks.

The presentation attack types are diverse, which vary from typical printed facial photos and video replays to more high-priced 3D masks. Attackers can also utilize spoof attacks that the system designers are unaware of; we call them unknown spoof attack types. As the increasing growth of presentation attacks, face anti-spoofing systems have high probability of encountering unseen presentation instruments.

CNN-based methods can detect pre-defined attack types well, while the generalization ability of these methods is limited for the seen attack types in the training data. The

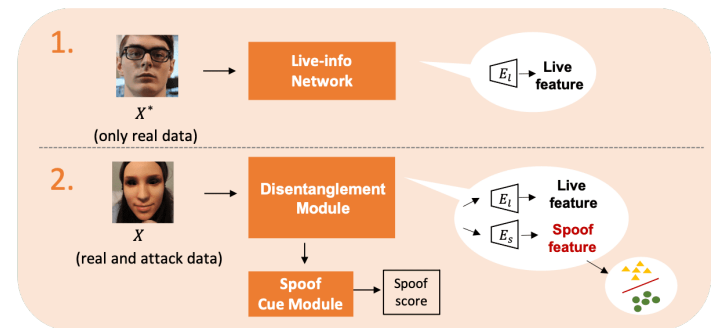


Figure 1: The pipeline of the proposed Dual-stage Feature Learning FAS system. The live-info network concentrates on learning live features in 1<sup>st</sup> stage. In the 2<sup>nd</sup> stage, the disentanglement module adopts the pre-trained Live-encoder in the 1<sup>st</sup> stage as a fixed live feature extractor and further disentangle spoof features and live features.

model trained with the existing spoof types would easily fail to detect unseen types of spoof attacks. Some CNN-based methods employ auxiliary information such as heart signal and facial depth to guide the network to distinguish real humans and spoof attacks [9, 17, 30, 14]. Although auxiliary supervision helps extract more discriminative features, the performances of FAS models trained with auxiliary supervisions rely on the accuracy of the additional tasks. Besides, existing auxiliary information may not be suited for all attack types, especially unseen attacks.

Detection of unseen attacks is a serious problem. It is impossible to define all auxiliary distinguishing information to guide the model training for detecting spoofs. The most crucial part lies in the model learning the essence of distinguishing real faces from spoof faces to avoid the network overfitting to training data. Disentangled representation learning is an effective approach to separate the model's latent representation into interpretable parts, further improving the model's robustness [33, 19]. These works proposed different strategies to enforce the model learning disentan-

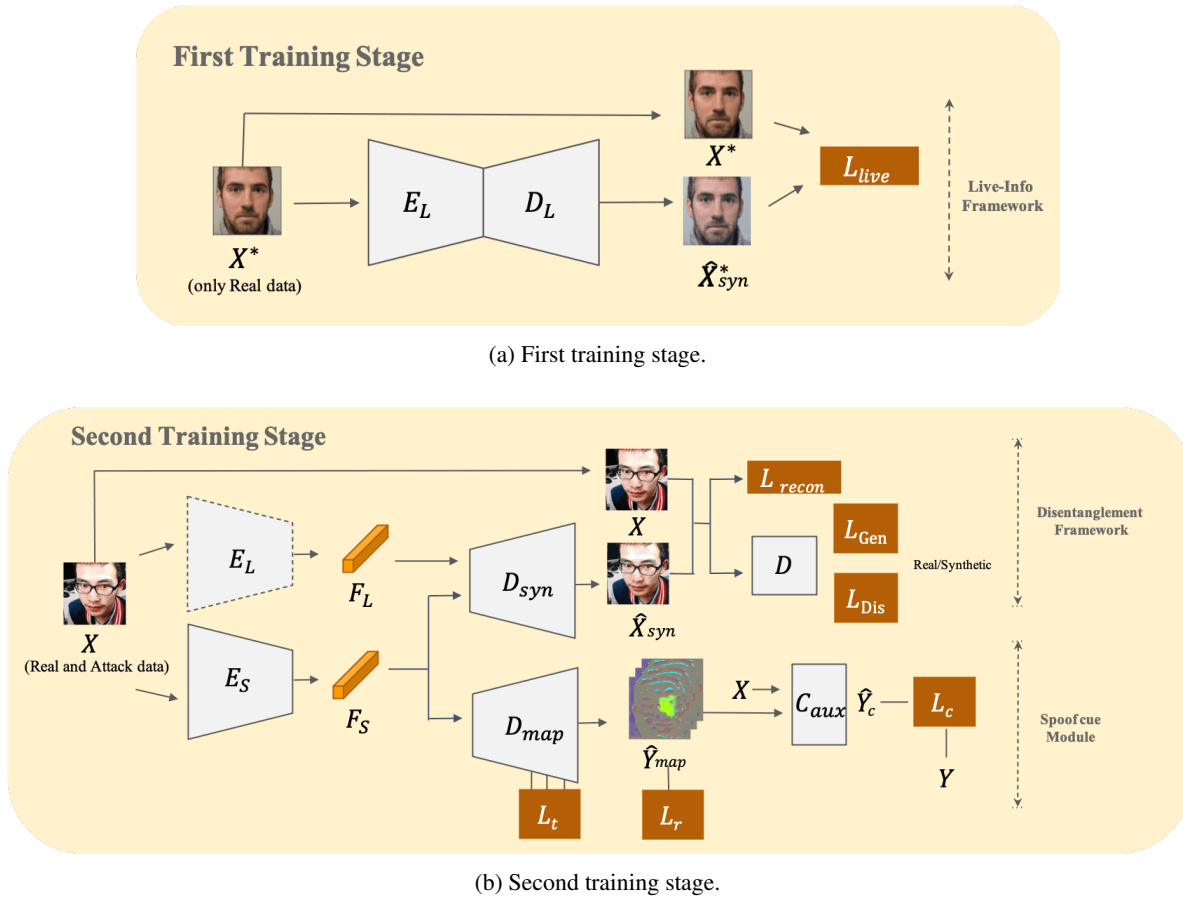


Figure 2: **Overview of the proposed dual-stage feature learning FAS system:** In the first stage(a), live-info network focuses on learning live features. In the second stage (b), we take pre-trained encoder  $E_L$  as live feature fixed extractor. By reconstructing original data, disentanglement module can concentrate on extracting disentangled spoof features.

gled features in an end-to-end architecture, and they use adversarial learning to facilitate disentangled representation. The stability of the generator and discriminator would largely influence the disentanglement representation effect. Furthermore, it is difficult to guarantee that the models can precisely extract disentangled representation for live and spoof components at the same time.

In light of these problems, we aim to improve the unknown attack detection accuracy with a two-stage disentanglement representation architecture to guarantee that the model extracts more stable disentangled representations. While many previous works leverage auxiliary supervision during training, we build our model without supervision from auxiliary tasks, whose ground truths may not be accurate or suitable for the target task.

The main contributions of our paper are summarized as follows:

1. We propose a novel dual-stage feature disentangled

representation learning method for the FAS problem. Our disentangled spoof features provide powerful generalization ability when facing unknown attacks.

2. We propose to regress spoof maps through our disentangled spoof feature. This method can further improve our model generalization ability against unknown attacks.
3. Our experiments show that our proposed method outperforms the state-of-the-art methods on several cross-type FAS protocols.

## 2. Related Work

### 2.1. Face Anti-Spoofing

Traditional FAS methods utilize classical local descriptors such as LBP [20], HOG [13], and SIFT [23] to extract hand-crafted features as the facial images textures. After that, they use classic classifiers like SVM and LDA to make

a binary decision. Other methods focus on motion-based face anti-spoofing approaches; they use motion cues like eye-blinking [21], mouth movement [5], and head movement [12] to detect print-out attacks. However, motion-based methods fail against 3D mask attacks since some masks have well-preserved motion cues like eyes and mouth [15].

With the development of CNN, many researchers introduce the CNN model as a feature extractor and formulate FAS as a binary classification problem [7, 31, 9]. In order to learn discriminative features efficiently, they adopt auxiliary supervisions like depth [17], reflection [10], and rRPG [17, 16]. Despite applying task-oriented knowledge, deep learning based approaches have limited generalization capability. They tend to overfit to pre-defined spoof attack types in the training data. When the model encounters unseen spoof attacks, the FAS performance would drop drastically.

## 2.2. Disentangled representation learning

Disentangled representation is an unsupervised learning method that can untangle data into interpretable components. The well-disentangled spoof-related features could help us enhance the robustness of the FAS model. GAN and VAE architectures have the powerful capability to generate new data. Therefore, researchers utilize them to learn disentangled representations [28, 8]. [19] disentangles spoof traces into a hierarchical combination and further synthesizes new data to enhance training. [33] separates representation into liveness and content features, and they combine depth and texture information to boost training. These FAS works focus on solving common FAS problems; therefore, they utilize end-to-end strategies to limit their model disentangling spoof features. However, to detect unknown spoofing attacks, we require more generalized spoofing features. Different from previous work, We propose a dual-stage feature-disentangled representation learning method to extract more discriminative spoofing features.

## 3. Proposed method

We assume that every spoof face is composed of its real counterpart and spoof-related information. Accordingly, we split the latent representation of facial data into live feature and spoof feature. Live features correlate with the information that spoof images share with their real counterparts, such as identification and background. Spoof features are the spoof-related information that differentiates between attack-face and real-face images.

To disentangle the live feature and spoof feature efficiently, we propose a dual-stage feature learning method to disentangle face representation. Our training process has two stages. The first stage is based on the live-info network, and the second stage consists of the disentanglement mod-

ule and Spoof Cue module. The whole learning process is illustrated in Figure 2.

### 3.1. Live-info Network

In the first training stage, we only utilize real images as training data to train the live-info network. Figure 2a shows the network architecture of our live-info network. The Live-encoder  $E_L$  is a one-class auto-encoder. We use loss  $L_{Live}$  to constrain the model to extract live features and reconstruct real face images accurately. Because the model only learns information from real samples, it can extract the universal factors of variation from real data. We regard all the features extracted by the Live-encoder  $E_L$  as live features.

Let  $X^*$  denote the real training data. The live-info network can generate output  $\hat{X}_{syn}^* = D_L(E_L(X^*))$ . The loss function  $L_{Live}$  is given by

$$L_{Live} = \mathbb{E}_{x \sim P_X} [\|x^* - \hat{x}_{syn}^*\|_2^2] \quad (1)$$

### 3.2. Disentanglement Module

As shown in Figure 2b, we utilize both real and attack samples as training data in the second stage. Two-branch encoders  $E_L$  and  $E_S$  are used to extract the live features  $F_L$  and spoof features  $F_S$ , respectively. The Live-encoder  $E_L$  loads the pre-trained model weights in the first stage as a live feature fixed extractor. Since the Live-encoder  $E_L$  had not learned any information about spoof images, it can focus on extracting spoof-invariant features. In order to reconstruct complete input data, the Spoof-encoder  $E_S$  concentrates on extracting the information that differentiates attack and real data. We employ the loss  $L_{recon}$  and adversarial loss to facilitate the outputs  $\hat{X}_{syn}$  similar to the original images  $X$ . We take the generated data  $\hat{X}_{syn}$  and original data  $X$  as the input of the discriminator  $D$ , and it will determine if the input samples are real or synthetic.

Let  $(X, Y)$  denote the input data and the corresponding labels from all the training data. The adversarial loss is formed by

$$L_{Gen} = \mathbb{E}_{x \sim P_X} [(D(\hat{x}_{syn}) - 1)^2] \quad (2)$$

and

$$L_{Dis} = \mathbb{E}_{x \sim P_X} [(D(x) - 1)^2] + \mathbb{E}_{x \sim P_X} [(D(\hat{x}_{syn}))^2] \quad (3)$$

The reconstruction loss  $L_{recon}$  is given by

$$L_{recon} = \mathbb{E}_{x \sim P_X} [\|x - \hat{x}_{syn}\|_2^2] \quad (4)$$

### 3.3. Spoof Cue Module

We use the disentangled spoof feature to decode spoof maps in an unsupervised way. It can avoid our approach

from overfitting to known attacks in the training data. We modify the method proposed in [6] as our Spoof Cue Module. They adopt a U-Net[25] to generate spoof cue maps. In our Spoof Cue Module, we utilize disentangled spoof features  $F_S$  in the disentanglement module as input and adopt a decoder  $D_{map}$  to generate spoof maps.

The loss function  $L_t$  is used to facilitate contracting the spoof features  $F_S$  of real face images and separating the spoof features between real and spoof data. The loss  $L_t$  is applied in the last three layers of decoder  $D_{map}$ . Let  $f_i^a, f_i^p, f_i^n$  denote the feature vectors of anchor (real), positive (real), and negative (attack) samples, respectively, for the  $i$ -th triplet. We always pick live sample as the anchor. The loss function  $L_t$  consists of two parts, i.e.,  $L_t^{normal}$  and  $L_t^{hard}$ .

The loss  $L_t^{normal}$  counts normal triplet loss for all valid triplets and averages those whose values are positive.  $N$  denotes the number of triplets.  $\alpha$  is a pre-defined margin constant.

$$L_t^{normal} = \frac{1}{N} \sum_{i=1}^N \max(\|f_i^a - f_i^p\|_2^2 - \|f_i^a - f_i^n\|_2^2 + \alpha, 0) \quad (5)$$

The loss  $L_t^{hard}$  is a hard-triplet loss. For each anchor  $f_i^a$ , we take the pair  $(f_i^a, f_j^p)$  which has the largest Euclidean distance as the hard positive pair. The positive sample index  $j$  is selected from  $L$  and  $j \neq i$ , where  $L$  denotes the set of all live samples in the current batch. For each anchor  $f_i^a$ , we find the pair  $(f_i^a, f_k^n)$  with the smallest Euclidean distance as the hard negative pair. The spoofing sample index  $k$  is selected from  $N$ , where  $N$  denotes the set of all spoofing samples in the current batch.  $T$  denotes the number of triplets.  $m$  is the pre-defined margin constants. For each triplet  $(f_i^a, f_j^p, f_k^n)$ , we compute each hard triplet loss, and then average them.

$$L_t^{hard} = \frac{1}{T} \sum_{i=1}^T \max(\max_{j \in L} \|f_i^a - f_j^p\|_2^2 - \min_{k \in N} \|f_i^a - f_k^n\|_2^2 + m, 0) \quad (6)$$

where  $m$  is a pre-defined margin.

The loss function  $L_t$  is the sum of the above two loss functions, i.e.

$$L_t = L_t^{normal} + L_t^{hard} \quad (7)$$

We need to generate the discriminative spoof map  $\hat{y}_{map}$ , which represents the difference between real data and attack data. For this purpose, we use  $L_r$  and  $L_c$  to facilitate training the decoder  $D_{map}$ .

We use loss  $L_r$  to minimize the spoof map of real data while without setting constraints for spoofing data.  $x \in P_{live}$  denotes the real input data.

$$L_r = \mathbb{E}_{x \sim P_{live}} [\|\hat{y}_{map}\|_1] \quad (8)$$

$C_{aux}$  is a binary classifier to strengthen the  $D_{map}$  decoding spoof map  $\hat{y}_{map}$ . We overlap spoof map  $\hat{y}_{map}$  with original data  $X$  as the input of  $C_{aux}$ . The loss  $L_c$  for training the classifier  $C_{aux}$  is given by

$$L_c = \mathbb{E}_{x \sim P_X, y \sim P_Y} -y \cdot \log(\hat{y}_c) + (1-y) \cdot \log(1-\hat{y}_c) \quad (9)$$

where  $y$  denotes the binary label, and  $\hat{y}_c$  is the prediction of the classifier  $C_{aux}$ .

### 3.4. Training and Testing

We weight the summation of the loss functions described above as the final loss in the second training stage, given by

$$L = \lambda_1 L_{recon} + \lambda_2 L_{Gen} + \lambda_3 L_t + \lambda_4 L_r + \lambda_5 L_c \quad (10)$$

where  $\lambda_1, \lambda_2, \lambda_3, \lambda_4$  and  $\lambda_5$  are the weights associated with the loss functions described above, and their values are set to 4, 1, 3, 5, and 5, respectively, in our experiments.

At the testing phase, we calculate the average of the generated spoof map  $\hat{y}_{map}$  as the spoof score and further utilize it to make the real/attack decision. The input data with large scores are likely to be a spoof one. We only need to use the encoder  $E_s$  and the decoder  $D_{map}$  in the inference stage; therefore, our method can achieve  $121.48 \pm 1.4$  FPS on GeForce GTX 1080.

## 4. Experiments

### 4.1. Databases and Protocols

We evaluate the proposed method on five popular FAS databases: SiW-M[18], MSU-MFSD[27], CASIA-MFSD[34], Replay-Attack[3], and SiW[17]. To evaluate the generalization ability of the proposed method in unseen attacks, we test on several cross-type protocols[1] [18] like SiW-M[18], MSU-MFSD[27], CASIA-MFSD[34], and Replay-Attack[3]. We also evaluate the intra-type protocol, which adopts identical attack types but with illumination and expression variations, like sub-protocol 1 and 2 in SiW dataset[17]. Moreover, we provide the cross-database result in the supplementary material.

### 4.2. Evaluation metrics

We select five metrics which are commonly used to evaluate the performance in FAS tasks: Equal Error Rate (EER), Attack Presentation Classification Error Rate (APCER) [2],

Method	Metrics(%)	Replay	Print	Mask attacks				Makeup Attacks			Partial Attacks			Average	
				Half	Silicone	Trans.	Paper	Manne.	Obfusc.	Imperson.	Cosmetic	Funnyeye	Paperglasses		Partialpaper
Auxiliary [17] (CVPR 2018)	APCER	23.7	7.3	27.7	18.2	97.8	8.3	16.2	100.0	18.0	16.3	91.8	72.2	0.4	38.3±37.4
	BPCER	10.1	6.5	10.9	11.6	6.2	7.8	9.3	11.6	9.3	7.1	6.2	8.8	10.3	8.9±2.0
	ACER	16.8	6.9	19.3	14.9	52.1	8.0	12.8	55.8	13.7	11.7	49.0	40.5	5.3	23.6±18.5
	EER	14.0	4.3	11.6	12.4	24.6	7.8	10.0	72.3	10.1	9.4	21.4	18.6	4.0	17.0±17.7
DTN [18] (CVPR 2019)	APCER	1.0	0.0	0.7	24.5	58.6	0.5	3.8	73.2	13.2	12.4	17.0	17.0	0.2	17.1±23.3
	BPCER	18.6	11.9	29.3	12.8	13.4	8.5	23.0	11.5	9.6	16.0	21.5	22.6	16.8	16.6±6.2
	ACER	9.8	6.0	15.0	18.7	36.0	4.5	7.7	48.1	11.4	14.2	19.3	19.8	8.5	16.8±11.1
	EER	10.0	<b>2.1</b>	14.4	18.6	26.5	5.7	9.6	50.2	10.1	13.2	19.8	20.5	8.8	16.1±12.2
DST [19] (ECCV 2020)	APCER	1.6	0.0	0.5	7.2	9.7	0.5	0.0	96.1	0.0	21.8	14.4	6.5	0.0	12.2±26.1
	BPCER	14.0	14.6	13.6	18.6	18.1	8.1	13.4	10.3	9.2	17.2	27.0	35.5	11.2	16.2±7.6
	ACER	<b>7.8</b>	7.3	7.1	12.9	<b>13.9</b>	4.3	6.7	53.2	4.6	19.5	20.7	21.0	5.6	14.2±13.2
	EER	<b>7.6</b>	3.8	8.4	13.8	14.5	5.3	4.4	35.4	0.0	19.3	21	20.8	1.6	12.0±10.0
HMP [30] (ECCV 2020)	APCER	12.4	5.2	8.3	9.7	13.6	0.0	2.5	30.4	0.0	12.0	22.6	15.9	1.2	10.3±9.1
	BPCER	13.2	6.2	13.1	10.8	16.3	3.9	2.3	34.1	1.6	13.9	23.2	17.1	2.3	12.2±9.4
	ACER	12.8	5.7	10.7	10.3	14.9	<b>1.9</b>	<b>2.4</b>	32.3	0.8	12.9	22.9	16.5	<b>1.7</b>	11.2±9.2
	EER	13.4	5.2	8.3	9.7	13.6	5.8	2.5	33.8	<b>0.0</b>	14.0	23.3	16.6	<b>1.2</b>	11.3±9.5
NAS-FAS [32] (TPAMI 2020)	APCER	12.8	9.0	9.7	13.1	19.1	1.1	5.4	31.0	0.0	15.0	15.1	18.6	5.0	11.9±8.4
	BPCER	10.1	6.8	13.1	11.1	12.5	2.8	0.0	26.1	0.8	15.3	17.8	13.5	2.3	10.2±7.5
	ACER	9.3	7.9	11.4	12.1	15.8	<b>1.9</b>	2.7	28.5	<b>0.4</b>	15.1	<b>16.5</b>	<b>16.0</b>	3.8	10.9±7.8
	EER	9.3	6.8	9.7	11.1	<b>12.5</b>	2.7	<b>0.0</b>	26.1	<b>0.0</b>	15.0	<b>15.1</b>	<b>13.4</b>	2.3	9.5±7.4
<b>Ours</b>	APCER	8.16	1.7	0.0	3.7	11.45	0.0	0.0	17.39	1.63	8.0	41.87	27.55	6.97	9.88±12.5
	BPCER	7.63	9.16	6.87	8.39	20.45	7.63	6.87	14.5	4.58	9.9	6.87	13.74	6.1	9.44±4.36
	ACER	7.89	<b>5.43</b>	<b>3.43</b>	<b>6.05</b>	15.9	3.81	3.43	<b>15.94</b>	3.1	<b>8.96</b>	24.37	20.64	6.54	<b>9.65±7.15</b>
	EER	7.63	6.87	<b>5.34</b>	<b>7.63</b>	12.97	<b>0.0</b>	6.1	<b>14.5</b>	3.05	<b>8.39</b>	16.03	17.55	6.1	<b>8.63±5.18</b>

Table 1: Comparison of the state-of-the-art methods on the cross-type testing on SiW-M dataset [18].

Bona Fide Presentation Classification Error Rate (BPCER) [2], the average of APCER and BPCER, Average Classification Error Rate (ACER) = (APCER+BPCER)/2 [2], and Area Under Curve (AUC).

### 4.3. Implementation Details

We crop the face regions from the datasets as training data by the face detector in Dlib [11] for the datasets without providing face location ground truth. Then, the cropped regions are all resized to the size 256x256. We resample the training data to retain the ratio between real images and attack images to 1:1. We choose Adam [4] as the optimizer with an initial learning rate of  $5e-4$ , and the training batch size is 32. The method is implemented by Pytorch[22]. Because our model is a two-stage training procedure, we train the live-info network with 10 epochs in the first stage. In the second stage, we load the pre-trained Live-encoder  $E_L$  as a fixed live-feature extractor to further train the disentanglement module. We set the switching frequency of the discriminator and disentanglement module empirically. Our disentanglement module can converge after around ten epochs for all the datasets in our implementation.

### 4.4. Testing on SiW-M

We evaluate the cross-type detection performance on the SiW-M dataset. The SiW-M dataset adopts the leave-one-out testing protocols and has 968 videos of 13 types of spoof attacks. Following the original work [18], we use 12 types of spoof attacks plus 80% of the real videos to train the model and evaluate the remaining one attack type plus the

20% of real videos in every sub-protocol. For a fair comparison with other state-of-the-art methods, we adopt the same approach describing in [18]. The threshold are decided with the training data and fix it for all testing sub-protocols. As Table 1 shows, our method achieves top performance with average ACER (9.65%) and EER (8.63%) for all the 13 attack types. Furthermore, our evaluation result has a smaller standard deviation than all the other methods. It shows that our method has a stronger generalization ability when encountering unknown attack types.

### 4.5. Testing on MSU-MFSD, CASIA-MFSD, and Replay-Attack

We also test the cross-type detection performance on MSU-MFSD, CASIA-MFSD, and Replay-Attack; we follow the unseen attack type protocol proposed in [1]. This protocol also adopts ‘leave one attack type out’ to validate the model robustness against unknown attack types. As illustrated in Table 2, our method exceeds the state-of-the-art results and achieves the best average AUC(98.0%) in this protocol. Specifically, we greatly improve the AUC (%) for the printed spoof type in MSU-MFSD, which type other methods perform worse. It demonstrates that our method is more robust when encountering unseen attacks.

### 4.6. Testing on SiW

The SiW dataset contains 4,478 videos. It defines three sub-protocols to evaluate the model generalization ability under different face expressions, mediums, and attack types, respectively. The evaluation results are reported in

Method	CASIA-MFSD			Replay-Attack			MSU-MFSD			Overall
	Video	Cut photo	Wrapped	Video	Digital Photo	Printed	Printed	HR video	Mobile Video	
SVM+LBP [35]	91.94	91.7	84.47	99.08	98.17	87.28	47.68	99.5	97.61	88.55±16.25
NN+LBP [29]	94.16	88.39	79.85	99.75	95.17	78.86	50.57	99.93	93.54	86.69±16.25
DTN [18]	90	97.3	97.5	99.9	99.9	99.6	81.6	99.9	97.5	95.9±6.2
AIM-FAS [24]	93.6	99.7	99.1	99.8	99.9	99.8	76.3	99.9	99.1	96.4±7.8
NAS-FAS [32]	99.62	100	100	99.99	99.89	99.98	74.62	100	99.98	97.1±8.9
<b>Ours</b>	98.05	99.36	99.83	99.59	99.96	100	86.78	100	98.39	<b>98.0±4.26</b>

Table 2: AUC (%) of the state-of-the-art methods for the cross-type testing on CASIA-MFSD[34], Replay-Attack[3], and MSU-MFSD[34].

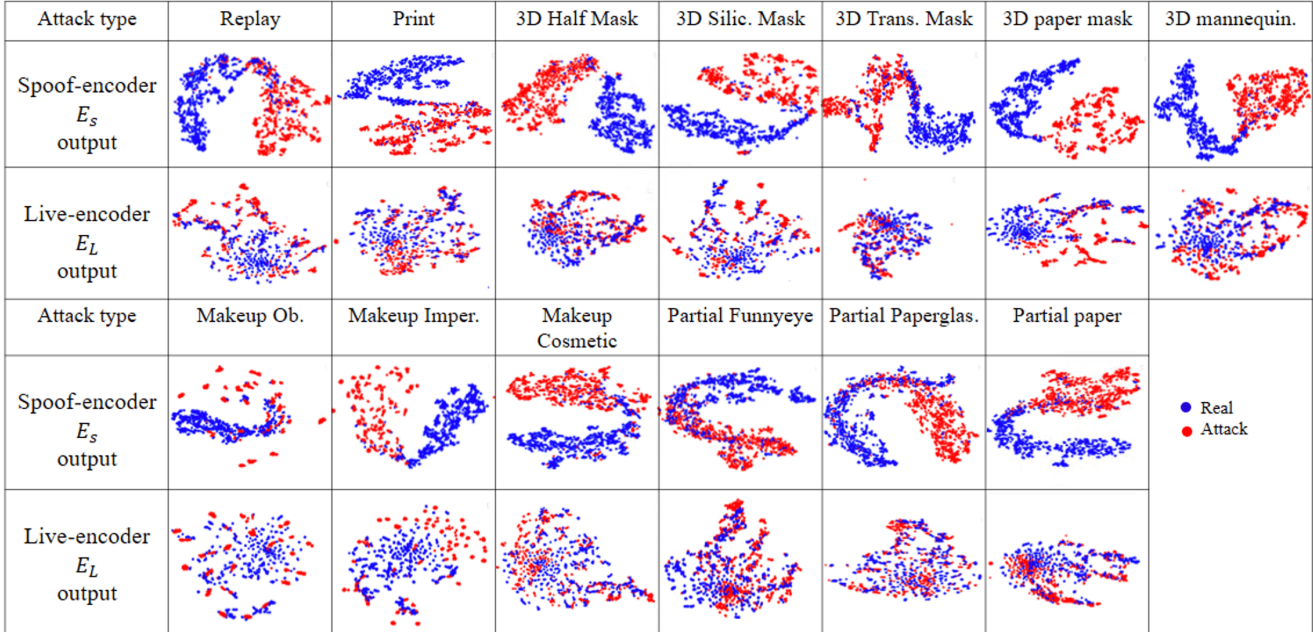


Figure 3: **Visualization of feature distributions of the SiW-M dataset by t-SNE[26].** We visualize the live feature  $F_L$  and spoof feature  $F_S$  in the SiW-M dataset of different unknown attack types.  $F_L$  of real and attack samples are highly overlapped. By contrast,  $F_S$  distinguishing.

Table 3. Our method performs the best in sub-protocol 1 and sub-protocol 2. In SiW protocol 3, the training data only contains one type of attack, which could possibly inhibit the disentangled feature learning in our network. We obtain competitive performance with 3.58 ACER (%). The experimental results show that our method has great generalization ability not only on the unseen attacks but also on different poses, attack mediums, and camera sensors.

## 5. Ablation Study

### 5.1. Discussion of the disentangled features

Our disentanglement module is designed to disentangle between spoof features and live features. To verify the rationality of the disentangled features, we swap the spoof features between a real image and an attack image to syn-

thesize a new translation image in this section. We choose two unpaired images A and B from the SiW dataset and the SiW-M dataset, where A is a real image, and B is an attack image. We experiment on the replay attack type because their screen textures can help us observe the discrepancy between real and attack images.

We use Live-encoder  $E_L$  and Spoof-encoder  $E_S$  to encode inputs into live features and spoof features, respectively. We obtain  $F_L^A$ ,  $F_S^A$ ,  $F_L^B$ , and  $F_S^B$  after we encode images A and B. We exchange the spoof feature  $F_S$  of A and B. Then, we use the decoder  $D_{syn}$  to synthesize new real data  $B^T$  and attack data  $A^T$ . This translation proce-

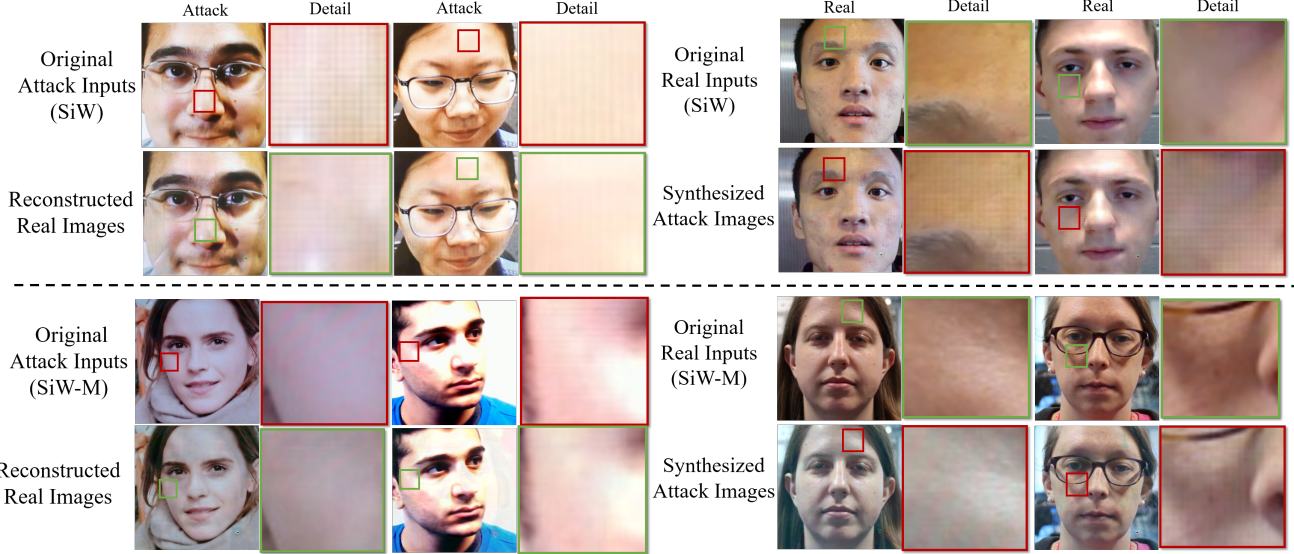


Figure 4: **Illustrations of the translation result in the SiW and the SiW-M datasets.** We swap the spoof features between a real person and a replay attack image to synthesize new translation images. The detailed images show that the screen textures have an apparent difference between the original inputs and the translation images.

Protocol	Method	APCER(%)	BPCER(%)	ACER(%)
1	HMP [30]	0.55	0.17	0.36
	DRL [33]	0.07	0.5	0.28
	DST [19]	0.00	0.00	<b>0.00</b>
	LGC [6]	0.00	0.5	0.25
	<b>Ours</b>	0.00	0.00	<b>0.00</b>
2	HMP [30]	0.08±0.17	0.15±0.00	0.11±0.08
	DRL [33]	0.08±0.17	0.21±0.16	0.15±0.14
	DST [19]	0.00±0.00	0.00±0.00	<b>0.00±0.00</b>
	LGC[6]	0.00±0.00	0.00±0.00	<b>0.00±0.00</b>
	<b>Ours</b>	0.00±0.00	0.00±0.00	<b>0.00±0.00</b>
3	HMP [30]	2.55±0.89	2.34±0.47	2.45±0.68
	DRL [33]	9.35±6.14	1.84±2.60	5.59±4.37
	DST [19]	8.3±3.3	7.3±3.3	7.9±3.3
	LGC [6]	1.61±1.69	0.77±1.09	<b>1.19±1.39</b>
	<b>Ours</b>	4.77±5.04	2.44±2.74	3.58±3.93

Table 3: Comparison of state-of-the-art methods on the three SiW protocols [17].

ture can be described as:

$$\begin{aligned}
 F_L^A &= E_L(A), F_S^A = E_S(A) \\
 F_L^B &= E_L(B), F_S^B = E_S(B) \\
 A^T &= D_{syn}(F_L^A, F_S^B), B^T = D_{syn}(F_L^B, F_S^A)
 \end{aligned}
 \quad (11)$$

We show the translation results of some samples from SiW and SiW-M datasets in Figure 4. Images in the first and the third row are original inputs. The second and the fourth row depicts the generated outputs by the above translation process. We use red boxes to point out the details of the attack images, and green boxes refer to the details of the

real images. As shown in the figure, we can easily observe the notable local detailed difference between original images and translation results. The original attack images are usually with repetitive striped patterns. After the translation process, we can notice that the translated images have fewer streaks on their surfaces. We also conduct the translation process for the original real data and generate the new synthetic attack images. We can notice that these new synthetic attack images contain some repetitive striped patterns. According to the translation results, we can confirm that our disentangled spoof features are critical for generating spoof images. With these spoof features, we can further employ these spoof information to distinguish attack data.

## 5.2. Contribution of each Loss function

In this section, we train the network with some components disabled or some loss excluded to compare the corresponding accuracies of the model. As shown in Table 4. The first-row method means that we train the method with the weights of the encoder  $E_L$  randomly initialized rather than use the pre-trained encoder  $E_L$  in the first stage. By comparing the results, we can see that *w/o* first stage has much lower accuracy in the experimental results. Our two-stage training can enhance the model performance by extracting more refined disentangled spoof features.

Then we train the model without using adversarial loss. We can note that the performance drops significantly in some attack types; it may be because the decoder is required to reconstruct the images only by  $L_{recon}$  loss. And then, we remove the  $C_{aux}$  classifier. Without using binary supervi-

Method	CASIA-MFSD			Replay-attack			MSU-MFSD			Overall
	Video	Cut photo	Wrapped	Video	Digital Photo	Printed	Printed	HR video	Mobile Video	
Proposed w/o first-stage	96.25	96.88	98.8	84.53	99.76	100	83.21	98.57	97.10	95.01±6.45
Proposed w/o discriminator	95.36	96.97	98.5	82.59	99.7	100	76.57	100	96.35	94.0±8.47
Proposed w/o aux classifier	96.69	98.47	97.77	93.65	99.23	99.92	76.5	82.07	95.42	93.3±8.29
Proposed w/o triplet loss	94.50	95.9	93.88	92.87	99.4	99.93	82.07	99.0	97.54	95.01±5.48
Proposed w normal triplet loss	96.66	98.83	99.72	97.34	99.76	100	84.25	99.4	99.28	97.25±5.01
<b>Proposed (full method)</b>	98.05	99.36	99.83	99.59	99.96	100	86.78	100	98.39	<b>98.0±4.26</b>

Table 4: **The ablation study of the proposed method.** Experimental results of removing some components from the proposed system with ‘leave one attack type out’ protocol in the CASIA-MFSD[34] dataset, the Replay-Attack[3] dataset, and the MSU-MFSD[34] dataset.

sion, the model is likely to overfit to the training set.

We also analyze the effectiveness of our loss  $L_t$ . One is only using normal triplet loss, and the other is abandoning the loss  $L_t$ . The poor accuracy shows that triplet learning can help the model decode more discriminative spoof maps. Using combinations of hard and normal triplet loss achieves better performance than merely using normal triplet. It indicates that the combinations of hard and normal triplet loss can facilitate contracting real samples and separating real samples and attack samples in the spoof feature space. Eventually, the proposed model training with all loss functions performs the best.

### 5.3. Visualization and Analysis

We visualize the feature embeddings  $F_L$  and  $F_S$  of our Live-encoder  $E_L$  and Spoof-encoder  $E_S$  with t-SNE [26] to justify our disentanglement representation strategy. We adopt the leave-one-out method on the SiW-M dataset of different unknown attack types. We randomly choose 1000 real data and 1000 unseen spoof-type data from the testing set. The live features represent the spoof-irrelevant feature in both attack samples and real samples like background and identification; hence, the live feature distribution of real data and attack data overlap at a high degree. On the other hand, our Spoof encoder  $E_S$  focuses on extracting the discrepancy between attacks and real samples. Therefore, the spoof feature visualization result of real data and attack data are more separated. The input images of the Spoof-encoder  $E_S$  and Live-encoder  $E_L$  are the same, but their visualization results are highly dissimilar. The difference between their visualization results demonstrates that our method can extract disentangled features. Additionally, we show that the feature embeddings  $F_L$  and  $F_S$  are disentangled in the supplementary material.

We also visualize the output of the decoder  $D_{map}$  to describe how we detect spoof attack. We test the model with unseen attack types, and the generated spoof maps are depicted in Figure 5. The first row shows the input images, and the second row depicts the generated spoof maps. As shown in Figure 5, the spoof maps for real face images are nearly zero maps, and there are some high-value attention

positions in the spoof maps for attack samples.




Figure 5: Examples of testing data in the SiW-M dataset and the corresponding spoof maps obtained by the proposed network.

## 6. Conclusion

We proposed a disentangled representation Network with dual-stage feature learning for face anti-spoofing. Unlike previous end-to-end architecture works, the dual-stage training design can improve the training stability and effectively encode the features to discriminate between real and attack. We further learn unsupervised spoof maps through our disentangled spoof features without using auxiliary information. These spoof features demonstrate great generalization ability when facing attacks of unknown spoof types. Extensive experiments on public FAS datasets demonstrate that the proposed method achieves state-of-the-art performance for unseen spoof types FAS tasks. Our method has promising performance in unknown attack protocols. However, there is some room to progress when the training data only contains few attack types like the SiW dataset sub-protocol 3. In the future, we plan to extend our dual-stage method to remove more spoof-irrelevant factors, further improving our model generalization ability.

## References

- [1] Shervin Rahimzadeh Arashloo, Josef Kittler, and William Christmas. An anomaly detection approach to face spoofing detection: A new formulation and evaluation protocol. *IEEE access*, 5:13868–13882, 2017.



- [2] ISO/IEC JTC1 SC37 Biometrics. Information technology–biometric presentation attack detection–part 3: testing and reporting. 2017.
- [3] Ivana Chingovska, André Anjos, and Sébastien Marcel. On the effectiveness of local binary patterns in face anti-spoofing. In *2012 BIOSIG-proceedings of the international conference of biometrics special interest group (BIOSIG)*, pages 1–7. IEEE, 2012.
- [4] Kingma Da. A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [5] Tiago de Freitas Pereira, Jukka Komulainen, André Anjos, José Mario De Martino, Abdenour Hadid, Matti Pietikäinen, and Sébastien Marcel. Face liveness detection using dynamic texture. *EURASIP Journal on Image and Video Processing*, 2014(1):2, 2014.
- [6] Haocheng Feng, Zhibin Hong, Haixiao Yue, Yang Chen, Keyao Wang, Junyu Han, Jingtuo Liu, and Errui Ding. Learning generalized spoof cues for face anti-spoofing. *arXiv preprint arXiv:2005.03922*, 2020.
- [7] Anjith George and Sébastien Marcel. Deep pixel-wise binary supervision for face presentation attack detection. In *2019 International Conference on Biometrics (ICB)*, pages 1–8. IEEE, 2019.
- [8] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 172–189, 2018.
- [9] Amin Jourabloo, Yaojie Liu, and Xiaoming Liu. Face de-spoofing: Anti-spoofing via noise modeling. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 290–306, 2018.
- [10] Taewook Kim, YongHyun Kim, Inhan Kim, and Daijin Kim. Basn: Enriching feature representation using bipartite auxiliary supervisions for face anti-spoofing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [11] Davis E King. Dlib-ml: A machine learning toolkit. *The Journal of Machine Learning Research*, 10:1755–1758, 2009.
- [12] Klaus Kollreider, Hartwig Fronthaler, Maycel Isaac Faraj, and Josef Bigun. Real-time face detection and motion analysis with application in “liveness” assessment. *IEEE Transactions on Information Forensics and Security*, 2(3):548–558, 2007.
- [13] Jukka Komulainen, Abdenour Hadid, and Matti Pietikäinen. Context based face anti-spoofing. In *2013 IEEE Sixth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–8. IEEE, 2013.
- [14] Bofan Lin, Xiaobai Li, Zitong Yu, and Guoying 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*, pages 61–68, 2019.
- [15] Siqi Liu, Baoyao Yang, Pong C Yuen, and Guoying Zhao. A 3d mask face anti-spoofing database with real world variations. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 100–106, 2016.
- [16] Si-Qi Liu, Xiangyuan Lan, and Pong C Yuen. Remote photoplethysmography correspondence feature for 3d mask face presentation attack detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 558–573, 2018.
- [17] Yaojie Liu, Amin Jourabloo, and Xiaoming Liu. Learning deep models for face anti-spoofing: Binary or auxiliary supervision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 389–398, 2018.
- [18] Yaojie Liu, Joel Stehouwer, Amin Jourabloo, and Xiaoming Liu. Deep tree learning for zero-shot face anti-spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4680–4689, 2019.
- [19] Yaojie Liu, Joel Stehouwer, and Xiaoming Liu. On disentangling spoof trace for generic face anti-spoofing. In *European Conference on Computer Vision*, pages 406–422. Springer, 2020.
- [20] Jukka Määttä, Abdenour Hadid, and Matti Pietikäinen. Face spoofing detection from single images using micro-texture analysis. In *2011 international joint conference on Biometrics (IJCB)*, pages 1–7. IEEE, 2011.
- [21] Gang Pan, Lin Sun, Zhaohui Wu, and Shihong Lao. Eyeblink-based anti-spoofing in face recognition from a generic webcam. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007.
- [22] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [23] Keyurkumar Patel, Hu Han, and Anil K Jain. Secure face unlock: Spoof detection on smartphones. *IEEE transactions on information forensics and security*, 11(10):2268–2283, 2016.
- [24] Yunxiao Qin, Chenxu Zhao, Xiangyu Zhu, Zezheng Wang, Zitong Yu, Tianyu Fu, Feng Zhou, Jingping Shi, and Zhen Lei. Learning meta model for zero-and few-shot face anti-spoofing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11916–11923, 2020.
- [25] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [26] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [27] Di Wen, Hu Han, and Anil K Jain. Face spoof detection with image distortion analysis. *IEEE Transactions on Information Forensics and Security*, 10(4):746–761, 2015.
- [28] Taihong Xiao, Jiapeng Hong, and Jinwen Ma. Elegant: Exchanging latent encodings with gan for transferring multiple face attributes. In *Proceedings of the European conference on computer vision (ECCV)*, pages 168–184, 2018.
- [29] Fei Xiong and Wael AbdAlmageed. Unknown presentation attack detection with face rgb images. In *2018 IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, pages 1–9. IEEE, 2018.

- [30] Zitong Yu, Xiaobai Li, Xuesong Niu, Jingang Shi, and Guoying Zhao. Face anti-spoofing with human material perception. In *European Conference on Computer Vision*, pages 557–575. Springer, 2020.
- [31] Zitong Yu, Yunxiao Qin, Xiaobai Li, Zezheng Wang, Chenxu Zhao, Zhen Lei, and Guoying Zhao. Multi-modal face anti-spoofing based on central difference networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 650–651, 2020.
- [32] Zitong Yu, Jun Wan, Yunxiao Qin, Xiaobai Li, Stan Z Li, and Guoying Zhao. Nas-fas: Static-dynamic central difference network search for face anti-spoofing. *arXiv preprint arXiv:2011.02062*, 2020.
- [33] Ke-Yue Zhang, Taiping Yao, Jian Zhang, Ying Tai, Shouhong Ding, Jilin Li, Feiyue Huang, Haichuan Song, and Lizhuang Ma. Face anti-spoofing via disentangled representation learning. In *European Conference on Computer Vision*, pages 641–657. Springer, 2020.
- [34] Zhiwei Zhang, Junjie Yan, Sifei Liu, Zhen Lei, Dong Yi, and Stan Z Li. A face antispoofing database with diverse attacks. In *2012 5th IAPR international conference on Biometrics (ICB)*, pages 26–31. IEEE, 2012.
- [35] Boulkenafet Zinelabidine, Komulainen Jukka, Lei Li, Xiao Feng, and Abdenour Hadid. Oulunpu: a mobile face presentation attack database with real-world variations. In *Proc. IEEE Int. Conf. on Identity, Security and Behavior Analysis, ISBA*, pages 1–7, 2017.