

# SLIP: Spoof-Aware One-Class Face Anti-Spoofing with Language Image Pretraining

Pei-Kai Huang<sup>1</sup>, Jun-Xiong Chong<sup>1</sup>, Cheng-Hsuan Chiang<sup>1</sup>, Tzu-Hsien Chen<sup>1</sup>, Tyng-Luh Liu<sup>2</sup>, Chiou-Ting Hsu<sup>1</sup>

<sup>1</sup> National Tsing Hua University, Taiwan

<sup>2</sup> Academia Sinica, Taiwan

{alwayswithme,jxchong,jimmy890718,gapp111062570,}@gapp.nthu.edu.tw,  
liutyng@iis.sinica.edu.tw, ctsu@cs.nthu.edu.tw

## Abstract

Face anti-spoofing (FAS) plays a pivotal role in ensuring the security and reliability of face recognition systems. With advancements in vision-language pretrained (VLP) models, recent two-class FAS techniques have leveraged the advantages of using VLP guidance, while this potential remains unexplored in one-class FAS methods. The one-class FAS focuses on learning intrinsic liveness features solely from live training images to differentiate between live and spoof faces. However, the lack of spoof training data can lead one-class FAS models to inadvertently incorporate domain information irrelevant to the live/spoof distinction (*e.g.*, facial content), causing performance degradation when tested with a new application domain. To address this issue, we propose a novel framework called Spoof-aware one-class face anti-spoofing with Language Image Pretraining (SLIP). Given that live faces should ideally not be obscured by any spoof-attack-related objects (*e.g.*, paper, or masks) and are assumed to yield zero spoof cue maps, we first propose an effective language-guided spoof cue map estimation to enhance one-class FAS models by simulating whether the underlying faces are covered by attack-related objects and generating corresponding nonzero spoof cue maps. Next, we introduce a novel prompt-driven liveness feature disentanglement to alleviate live/spoof-irrelevant domain variations by disentangling live/spoof-relevant and domain-dependent information. Finally, we design an effective augmentation strategy by fusing latent features from live images and spoof prompts to generate spoof-like image features and thus diversify latent spoof features to facilitate the learning of one-class FAS. Our extensive experiments and ablation studies support that SLIP consistently outperforms previous one-class FAS methods.

**Code** — <https://github.com/Pei-KaiHuang/AAAI25-SLIP>

## Introduction

Face anti-spoofing (FAS) plays a critical role in ensuring the security and reliability of face recognition systems, which may be deceived by facial spoof attacks, including print attacks (printing faces on paper/photos), replay attacks (displaying faces on electronic screens), or 3D mask attacks (wearing 3D masks). To counter potential facial spoof attacks, many two-class FAS methods proposed using both live

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

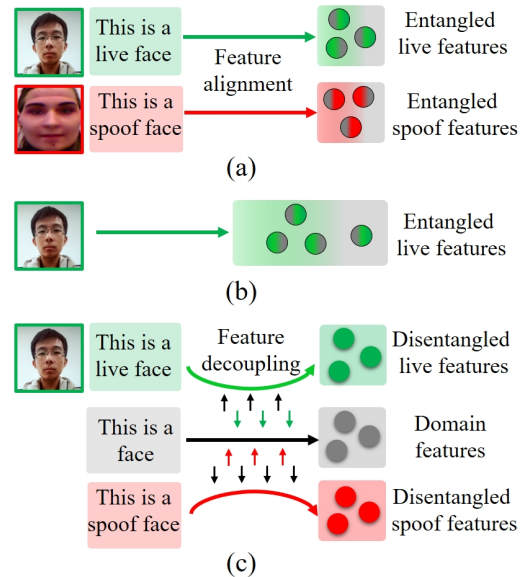


Figure 1: Liveness feature disentanglement: (a) Existing language-guided two-class FAS methods overlook the presence of domain information (*e.g.*, face content) in prompt learning to extract entangled live/spoof features. (b) One-class FAS methods learn entangled live features from live training images. (c) By exploring the mutual relatedness in the given text prompts, the proposed SLIP to one-class FAS learns pure live features via separating domain and live/spoof features.

and spoof training samples to learn FAS models to distinguish between live and spoof facial attacks. For example, in (Liu et al. 2024; Srivatsan, Naseer, and Nandakumar 2023; Fang et al. 2024), the authors proposed adopting contrastive language-image pretraining (CLIP) models (Radford et al. 2021) to align the live/spoof images with their corresponding prompts to learn FAS models. However, as shown in Figure 1 (a), these two-class FAS methods often disregard domain information existing in prompt learning. In particular, the term “face” is unrelated to “live” and “spoof”. As a result, these two-class methods easily overfit on seen attacks to learn entangled live/spoof features, which may cause performance

degradation when facing a new testing domain.

Since different live faces exhibit small distribution discrepancies between the training and test domains, some recent one-class FAS methods (Nikisins et al. 2018; Baweja et al. 2020; Lim et al. 2020; Huang, Xia, and Shen 2021; Huang et al. 2024a) proposed learning liveness information solely from live training samples to detect out-of-distribution (OOD) occurrences of spoof attacks. Although liveness information from live training samples alone is more feasible for identifying unseen spoof attacks from the OOD testing domain, one-class FAS models face a significant challenge posed by the absence of spoof training data. This challenge results in most previous one-class methods struggling to grasp the distinction between the “live” and “spoof” concepts for effectively distinguishing between live and spoof faces. In addition, because face anti-spoofing deals with highly similar visual characteristics (Huang et al. 2022b) between live and spoof faces (*e.g.*, the similar facial contents), one-class FAS methods are easier to involve live/spoof-irrelevance domain variations to learn entangled liveness features compared to two-class FAS methods, as shown in Figure 1 (b).

To address the above-mentioned challenges in one-class FAS, in this paper, we propose a novel framework called **Spoof-aware one-class face anti-spoofing with Language Image Pretraining (SLIP)**. First, to overcome the absence of spoof training data, we propose an effective language-guided spoof cue map estimation. Since live faces should ideally not be covered by any spoof-attack-related objects (*e.g.*, paper or masks) and are assumed to yield zero spoof cue maps (Huang et al. 2024a), we propose including prompt learning to simulate whether the underlying faces are covered by attack-related objects to produce corresponding nonzero or zero spoof cue maps. In particular, we create paired live/spoof prompts and corresponding spoof cue maps to guide one-class FAS models to grasp the distinction between “live” and “spoof” concepts. Next, we propose a novel prompt-driven liveness feature disentanglement to alleviate live/spoof-irrelative domain variations by disentangling live/spoof-relevant and domain-dependent information, as shown in Figure 1 (c). In particular, we create the content prompts to describe live/spoof-irrelative domain information (*e.g.*, face content), and then aggregate the latent content features extracted from the content prompts while distinguishing them from the latent live/spoof features extracted from the live/spoof prompts. Finally, we propose an effective spoof-like image feature augmentation by fusing latent features from live images and spoof prompts to generate spoof-like image features and thus diversify latent spoof features to facilitate the learning of one-class FAS. We conduct extensive experiments on seven public face anti-spoofing databases to evaluate the effectiveness of the proposed SLIP. Our experimental results on intra-domain and cross-domain testing demonstrate that the proposed SLIP surpasses existing one-class FAS approaches to achieve state-of-the-art performance.

Our contributions are summarized as follows:

- We develop a novel framework called **Spoof-aware one-class face anti-spoofing with Language Image Pretraining (SLIP)**. To the best of our knowledge, SLIP is the first

work focusing on learning disentangled liveness features for one-class face anti-spoofing.

- By adopting prompt learning to simulate whether live faces are covered by attack-related objects, the proposed language-guided spoof cue map estimation is able to guide one-class FAS models to grasp the distinction between “live” and “spoof” concepts, thereby distinguishing live images from spoof images.
- To alleviate live/spoof-irrelative domain variations, we propose a novel prompt-driven liveness feature disentanglement to learn disentangled liveness features from various prompts.
- Our extensive experimental results have shown that SLIP surpasses previous one-class FAS techniques to achieve state-of-the-art performance.

## Related Work

### One-Class Face Anti-Spoofing

One-class face anti-spoofing (FAS) focuses on learning intrinsic liveness features solely from live training images to differentiate between live and spoof faces (Huang et al. 2024b). In (Nikisins et al. 2018), the authors proposed using Gaussian Mixture Models (GMMs) and Image Quality Measure feature vectors proposed by (Wen, Han, and Jain 2015) to learn the distribution of live faces. Next, the authors in (Baweja et al. 2020) proposed generating pseudo-spoof features by mixing sampled Gaussian noise with live features to simulate out-of-distribution spoof features. Furthermore, the authors in (Lim et al. 2020) proposed reconstructing facial images to learn liveness features from live images. Finally, the authors in (Huang et al. 2024a) proposed adopting generative feature learning to yield latent spoof features for boosting the one-class FAS learning process.

### Two-Class Face Anti-Spoofing

Two-class FAS aims to learn discriminative liveness features from both live and spoof training images. For example, the authors in (Wang et al. 2024; Huang et al. 2022a, 2023b) proposed learning disentangled liveness features that are independent of domain-specific features to develop generalized FAS models. In (Zhou et al. 2024, 2022b,a), the authors proposed projecting, generating, or mixing diverse style information of potential spoof attacks to effectively mitigate domain gaps. The authors in (Yu et al. 2020b; Huang et al. 2022b, 2023a; Yu et al. 2021; Wang et al. 2020b; Chen et al. 2021) proposed using predefined or learnable descriptors to enhance the representation capability of standard convolution for effectively addressing seen spoofing attacks. To counter unseen types of spoof attacks, the authors in (Huang et al. 2023c) proposed employing contrastive learning to associate similar characteristics between seen and unseen spoofing attacks. In addition, the authors in (Liu et al. 2016a; Liu, Lan, and Yuen 2022, 2018, 2021; Yu et al. 2024; Huang, Chin, and Hsu 2021) proposed estimating remote photoplethysmography (rPPG) signals to distinguish live and spoof faces.

Some recent two-class face anti-spoofing methods focus on learning intrinsic liveness features from both live and spoof

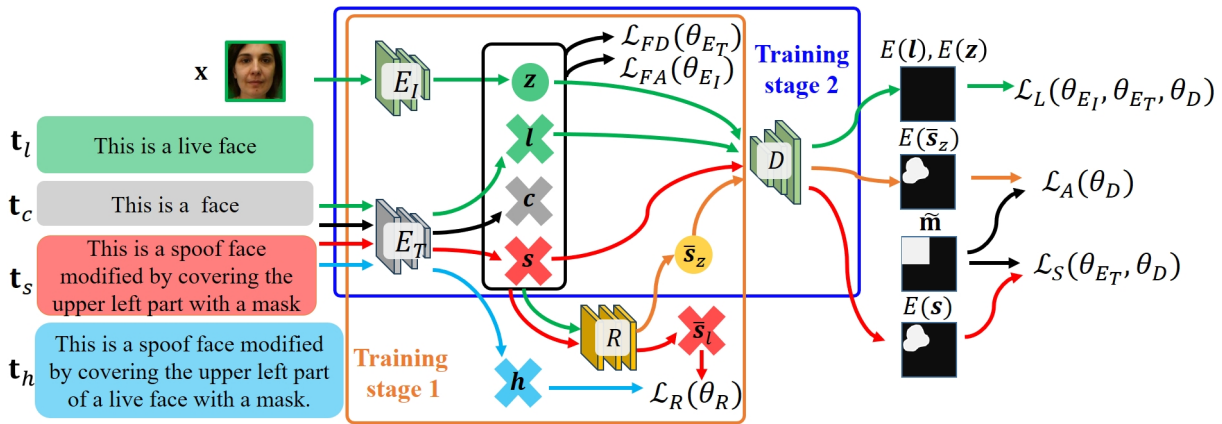


Figure 2: The proposed SLIP consists of one image encoder  $E_I$ , one text encoder  $E_T$ , one spoof cue map decoder  $D$ , and one fusion module  $R$ .

training images by using language guidance. In (Srivatsan, Naseer, and Nandakumar 2023), the authors first proposed incorporating language guidance to align features between image and text modalities, thereby enhancing the generalizability of face anti-spoofing (FAS) systems. Next, the authors in (Liu et al. 2024) proposed using the textual feature from large-scale vision-language models, *e.g.*, CLIP (Radford et al. 2021), to dynamically adjust the classifier’s weights for exploring generalizable visual features. Furthermore, in (Liu, Wang, and Yuen 2024), the authors proposed learning a domain-invariant prompt to address domain discrepancies among attack types for enhancing the performance of two-class FAS. In addition, the authors of (Fang et al. 2024) proposed including language modality into the FAS model to guide its attention mechanism to focus more effectively on the face region.

While language-guided strategies have shown considerable promise in enhancing two-class FAS methods, their potential remains unexplored in one-class FAS.

## Methodology

The problem of one-class face anti-spoofing (FAS) is inherently characterized by the absence of spoof facial images in the training data. This issue results in most previous one-class methods struggling to grasp the distinction between the “live” and “spoof” concepts. A recent technique by (Huang et al. 2024a) employs generative feature learning to yield latent spoof features, boosting the one-class FAS learning process. However, their method does not ensure that the generated spoof features align well with real-world spoofing scenarios, and consequently, may not consistently produce satisfactory results. (See our experiments.) Furthermore, live facial images encompass not only liveness but also domain properties (*e.g.*, facial content), as pointed out in previous research (Wang et al. 2020a, 2022). The subtlety is the primary reason that existing one-class FAS models tend to learn intertwined liveness features entangled with other domain-specific variations (*e.g.*, general facial features) rather than focusing narrowly on pure liveness attributes.

We illustrate the proposed SLIP framework in Figure 2. To leverage vision foundation models, SLIP employs the pre-trained CLIP (Radford et al. 2021) as the primary network backbone for the image encoder  $E_I$  and the text encoder  $E_T$ . To adapt CLIP for the one-class FAS task, we propose incorporating prompt learning to fine-tune  $E_I$  and  $E_T$  for learning distinct live/spoof features, thus effectively overcoming domain variations that are irrelevant to live/spoof classification. Next, we add a fusion module  $R$  that empowers our model to generate spoof-like image features by combining live image features and spoof prompt features. The introduction of  $R$  is crucial as it enriches the diversity of spoof features. Finally, we employ a spoof cue map decoder  $D$  to estimate the corresponding spoof cue maps (SCM) from the latent features to predict whether a face image is `Live` or `Spoof`. There are four distinct sets of parameters to be optimized during the model training:  $\theta_{E_I}$  of  $E_I$ ,  $\theta_{E_T}$  of  $E_T$ ,  $\theta_R$  of  $R$ , and  $\theta_D$  of  $D$ . We next elaborate on the details of our approach.

## Language-Guided Spoof Cue Map Estimation

In the task of one-class FAS, live faces should ideally not be obscured by any spoof-attack-related objects (*e.g.*, paper, or masks). In addition, a recent technique by (Huang et al. 2024a) has shown that live images should yield zero spoof cue maps (SCM), while spoof images should produce nonzero spoof cue maps. Therefore, to address the absence of spoof training data, we propose including prompt learning to enhance one-class FAS models by simulating whether the underlying faces are covered by attack-related objects and generating corresponding nonzero spoof cue maps.

**Various Prompts** In Table 1, we consider designing prompts that describe whether the facial region is obscured by spoof-attack-related objects as either live prompts  $t_l \in T_l$  or spoof prompts  $t_s \in T_s$ , where  $T_l$  and  $T_s$  are the sets of the live and spoof prompts, respectively. These prompts are used to assist one-class FAS models in understanding the concepts of “live” and “spoof”. Next, to simulate live/spoof-irrelative domain variations, we also design content prompts  $t_c \in T_c$  that do not include live/spoof-relevant information,

Live prompts $\mathbf{t}_l$	This is a {live / real / ...} face.
Spoof prompts $\mathbf{t}_s$	This is a {spooof / fake / ...} face modified by covering the {upper / right / ...} part with a {mask / photo / ...}.
Content prompts $\mathbf{t}_c$	This is a {face / facial / ...} image.
Hybrid prompts $\mathbf{t}_h$	This is a {spooof / fake / ...} modified by covering the {upper / right / ...} part of a {live / real / ...} face with a {mask / photo / ...}.

Table 1: Examples of various prompts.




Figure 3: Examples of  $\tilde{\mathbf{m}}$  with fixed size produced by using different positions as specified in the spoof prompts.

where  $T_c$  is the content prompt set. Finally, we include hybrid prompts  $\mathbf{t}_h \in T_h$  that simultaneously include live and spoof information, where  $T_h$  is the hybrid prompt set.

**Corresponding Pseudo SCMs** Observe from Table 1 that we include a specified placeholder to indicate the position of occluding objects. Hence, our next goal is to design the corresponding pseudo spoof cue maps  $\tilde{\mathbf{m}} \in \mathcal{M}$  for various prompts. First, we also adopt zero SCMs (*i.e.*,  $\mathbf{0}$ ) as the corresponding SCM for the live prompts. Next, based on the positions of occluding objects specified in the spoof prompts, we sample binary square masks with random sizes at corresponding positions to produce the nonzero SCMs for the spoof prompts. Figure 3 gives some basic samples of  $\tilde{\mathbf{m}}$  with fixed size produced by different positions as specified in  $\mathbf{t}_s$ . Because the different combinations of  $\tilde{\mathbf{m}} \in \mathcal{M}$  are able to cover the entire spatial region by specifying arbitrary objects to cover different facial regions within the spoof prompts, our method is capable of detecting potential spoof attacks in any facial region.

### Zero SCM Estimation from Live Images and Prompts

To learn the zero SCMs from live image features  $\mathbf{x} \in X$  and live prompts  $\mathbf{t}_l \in T_l$ , we define the liveness loss  $\mathcal{L}_L$  to constrain the image encoder  $E_I$ , the text encoder  $E_T$ , and the decoder  $D$  by,

$$\mathcal{L}_L = \mathcal{L}_I + \mathcal{L}_T, \quad (1)$$

which incorporates both the image liveness loss  $\mathcal{L}_I$  and the text liveness loss  $\mathcal{L}_T$  defined by

$$\mathcal{L}_I = \sum_{\mathbf{x} \in X} \|D(\mathbf{z}) - \mathbf{0}\|_2^2, \quad (2)$$

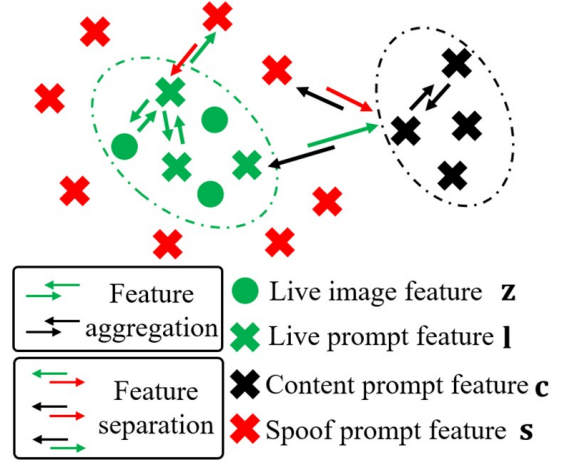


Figure 4: Illustration of liveness feature disentanglement.

and

$$\mathcal{L}_T = \sum_{\mathbf{t}_l \in T_l} \|D(\mathbf{l}) - \mathbf{0}\|_2^2, \quad (3)$$

where  $\mathbf{z} = E_I(\mathbf{x})$  and  $\mathbf{l} = E_T(\mathbf{t}_l)$  denote the live image features and the live prompt features, respectively.

### Nonzero SCM Map Estimation from Spoof Prompts

Now, with the paired spoof prompts  $\mathbf{t}_s$  and SCMs  $\tilde{\mathbf{m}}$ , the concern of grasping the distinction between the “live” and “spoof” concepts for one-class FAS models is no longer an issue. In particular, as shown in Figure 2, we define the spoof loss  $\mathcal{L}_S$  to constrain the text encoder  $E_T$  and the decoder  $D$  to predict the corresponding SCM  $\tilde{\mathbf{m}}$  by,

$$\mathcal{L}_S = \sum_{\mathbf{t}_s \in T_s, \tilde{\mathbf{m}} \in \mathcal{M}} \|D(\mathbf{s}) - \tilde{\mathbf{m}}\|_2^2, \quad (4)$$

where  $\mathbf{s}$  denotes the spoof prompt features.

### Prompt-Driven Liveness Feature Disentanglement

Observe from Table 1 that live prompts  $\mathbf{t}_l$  and spoof prompts  $\mathbf{t}_s$  simultaneously contain live/spoof-relevant information (*e.g.*, “live” or “covering with a mask”) and domain-dependent information (*e.g.*, “face”). The live and spoof features learned from  $\mathbf{t}_l$  and  $\mathbf{t}_s$  may be entangled with other domain-dependent variations. To address this issue, we propose to disentangle live/spoof-relevant and domain-dependent information to alleviate live/spoof-irrelative domain variations. In particular, as shown in Figure 4, we first define the feature disentanglement loss  $\mathcal{L}_{FD}$  to learn  $\theta_{E_T}$  of  $E_T$  from various prompts by

$$\mathcal{L}_{FD} = - \sum_{i=1}^{N_c} \sum_{j \neq i}^{N_c} \left( \log \frac{\exp(\cos(\mathbf{c}_i, \mathbf{c}_j))}{\sum_{p=1}^{N_l} \exp(\cos(\mathbf{c}_i, \mathbf{l}_p)) + \sum_{q=1}^{N_s} \exp(\cos(\mathbf{c}_i, \mathbf{s}_q))} \right) - \sum_{i=1}^{N_l} \sum_{j \neq i}^{N_l} \left( \log \frac{\exp(\cos(\mathbf{l}_i, \mathbf{l}_j))}{\sum_{k=1}^{N_s} \exp(\cos(\mathbf{l}_i, \mathbf{s}_k))} \right), \quad (5)$$

where  $N_l$ ,  $N_s$ , and  $N_c$  are the number of live prompts  $\mathbf{t}_l$ , spoof prompts  $\mathbf{t}_s$  and content prompts  $\mathbf{t}_c$ , respectively, and  $\mathbf{c} = E_T(\mathbf{t}_c)$  denotes the content prompt features. Note that, the denominators only include the negative pairs to remove the negative-positive-coupling effect (Yeh et al. 2022). In




Figure 5: Illustration of (a) spoof prompt feature reconstruction and (b) spoof-like image feature augmentation.

Equation 5, the first term is designed to learn disentangled liveness features by separating the latent live/spoof prompt features (*i.e.*,  $l$  and  $s$ ) and the latent content features (*i.e.*,  $c$ ). The second term is designed to aggregate the latent live features  $l$  while distinguishing them from the latent spoof features  $s$ .

With the disentangled liveness features learned with  $E_T$ , by aggregating the live prompt features  $l$  and the live image features  $z$ ,  $E_I$  is able to learn a disentangled feature space that reduces live/spoof-irrelative domain variations and focuses on learning live/spoof-relevant information, thereby effectively detecting the out-of-distribution (OOD) occurrences. Therefore, we define the feature alignment loss  $\mathcal{L}_{FA}$  to learn  $\theta_{E_I}$  of  $E_I$  by,

$$\mathcal{L}_{FA} = - \sum_{i=1}^{N_l} \sum_{j=1}^{N_z} \left( \log \frac{\exp(\cos(l_i, z_j))}{\sum_{k=1}^{N_s} \exp(\cos(l_i, s_k))} \right) \quad (6)$$

where  $N_z$  is the number of live images  $\mathbf{x} \in X$ .

### Spoof-Like Image Feature Augmentation

With the learned disentangled live and spoof features, our next goal is to diversify latent spoof features to facilitate the learning of one-class FAS. Hence, we consider fusing the disentangled live image features  $z$  and the disentangled spoof prompt features  $s$  to augment the spoof-like image features  $\bar{s}_z$ , as shown in Figure 5. In particular, in Figure 5 (a), we first extract latent features  $l$ ,  $s$ , and  $h$  from the paired prompts  $t_l$ ,  $t_s$ , and  $t_h$  (shown in Table 1) and define the reconstruction loss  $\mathcal{L}_R$  to guide the fusion module  $R$  to reconstruct the spoof prompt features  $\bar{s}_l$  by,

$$\begin{aligned} \mathcal{L}_R &= \sum_{t_l \in T_l, t_s \in T_s, t_h \in T_h} \|R(l, s) - \mathbf{h}\|_2^2 \quad (7) \\ &= \sum_{t_l \in T_l, t_s \in T_s, t_h \in T_h} \|\bar{s}_l - \mathbf{h}\|_2^2, \end{aligned}$$

where  $\mathbf{h} = E_T(t_h)$  denotes the spoof prompt features extracted from  $t_h$ . With the effect of  $\mathcal{L}_R$ , the model training would drive  $R$  to learn the fusion capability by combining the latent live features  $l$  and the latent spoof features  $s$  to approximate  $\mathbf{h}$ . As stated previously, we have aligned the live image features  $z$  and the live prompt features  $l$  in Equation (6). Thus, as shown in Figure 5 (b), it is reasonable to adopt the learned  $R$  to augment the spoof-like image features  $\bar{s}_z \in \bar{S}$  by fusing the live image features  $z$  and the spoof prompt features  $s$  by,

$$\bar{s}_z = R(z, s), \quad (8)$$

Type	Method	P	APCER	BPCER	ACER
1-class	IQM-GMM ( <i>ICB 18</i> )	1	75.35	18.56	46.95
	OC-fPAD ( <i>IJCB 20</i> )		38.63	21.85	30.24
	OC-LCFAS ( <i>Access 20</i> )		43.54	36.5	40.02
	AAE ( <i>CCBR 21</i> )		47.13	26.67	36.9
	OC-SCMNet ( <i>CVPR 24</i> )		20.83	26.15	23.49
	SLIP (Ours)		12.36	16.8	<b>14.58</b>
	IQM-GMM ( <i>ICB 18</i> )	2	41.56	27.78	34.67
	OC-fPAD ( <i>IJCB 20</i> )		51.81	19.83	35.82
	OC-LCFAS ( <i>Access 20</i> )		72.19	18.5	45.35
	AAE ( <i>CCBR 21</i> )		37.28	39.0	38.14
	OC-SCMNet ( <i>CVPR 24</i> )		22.05	28.81	25.43
	SLIP (Ours)		22.16	23.18	<b>22.67</b>
IQM-GMM ( <i>ICB 18</i> )	3	57.17±16.79	16.5±6.95	36.83±5.35	
OC-fPAD ( <i>IJCB 20</i> )		45.39±12.82	18.28±16.21	31.83±6.99	
OC-LCFAS ( <i>Access 20</i> )		38.51±13.08	39.52±11.13	39.02±2.16	
AAE ( <i>CCBR 21</i> )		26.62±13.67	52.93±16.09	39.77±3.74	
OC-SCMNet ( <i>CVPR 24</i> )		27.10±12.57	20.55±11.12	23.83±3.14	
SLIP (Ours)		26.35±9.76	20.03±5.87	<b>23.19±2.86</b>	
IQM-GMM ( <i>ICB 18</i> )	4	53.42±14.08	16.67±8.38	35.04±3.95	
OC-fPAD ( <i>IJCB 20</i> )		60.25±16.49	10.67±10.37	35.46±5.43	
OC-LCFAS ( <i>Access 20</i> )		36.91±10.24	20.5±8.01	28.07±5.32	
AAE ( <i>CCBR 21</i> )		26.33±18.5	40.17±29.04	33.12±8.9	
OC-SCMNet ( <i>CVPR 24</i> )		16.41±14.0	11.66±9.42	14.04±4.9	
SLIP (Ours)		15.02±3.84	10.9±4.66	<b>12.96±5.72</b>	

Table 2: Intra-domain testing on OULU-NPU.

where  $\bar{S} = \{\bar{s}_z\}$  is the augmented spoof feature set.

Next, we propose to estimate the SCMs from the augmented spoof-like image features  $\bar{s}_z$ . For each  $\bar{s}_z$ , as shown in Figure 5 (b), because the spoof prompt features  $s$  are extracted from spoof prompts  $t_s$ , the spoof cue map of  $\bar{s}_z$  should be the same as the corresponding spoof cue map  $\tilde{\mathbf{m}}$  of  $t_s$ . Therefore, in Figure 2, we define the augmented spoof loss  $\mathcal{L}_A$  to constrain  $D$  to learn the SCM estimation from  $\bar{s}_z$  by,

$$\begin{aligned} \mathcal{L}_A &= \sum_{\mathbf{x} \in X, t_s \in T_s, \tilde{\mathbf{m}} \in \mathcal{M}} \|D(R(z, s)) - \tilde{\mathbf{m}}\|_2^2 \quad (9) \\ &= \sum_{\mathbf{x} \in X, t_s \in T_s, \tilde{\mathbf{m}} \in \mathcal{M}} \|D(\bar{s}_z) - \tilde{\mathbf{m}}\|_2^2. \end{aligned}$$

Finally, by freezing the parameters  $\theta_R$  of  $R$ , the model training of the proposed SLIP can be achieved with

$$\begin{aligned} \theta_{E_I}^*, \theta_{E_T}^*, \theta_D^* &= \arg \min_{\theta_{E_I}, \theta_{E_T}, \theta_D} (\mathcal{L}_L(\theta_{E_I}, \theta_{E_T}, \theta_D) + \mathcal{L}_S(\theta_{E_T}, \theta_D) \\ &\quad + \lambda \mathcal{L}_{FD}(\theta_{E_T}) + \lambda \mathcal{L}_{FA}(\theta_{E_I}) + \mathcal{L}_A(\theta_D)), \quad (10) \end{aligned}$$

where  $\lambda = 0.8$  is weighted factors.

### Training and Testing

**Training** We iteratively optimize the four coupled optimization problems of (1), (4), (5), (6), (7), and (9) in an alternate manner. Before training stage, we first use (5) and (6) to learn better initialization weights for  $E_T$  and  $E_I$ , respectively, and then use (7) to learn better initialization weights for  $R$ . In each iteration, we first update  $E_T$ ,  $E_I$ , and  $R$  by minimizing  $\mathcal{L}_{FD}$  in (5),  $\mathcal{L}_{FA}$  in (6), and  $\mathcal{L}_R$  in (7) in the training stage 1, as shown in Figure 2. Next, in the training stage 2, we fix  $R$  and train  $E_T$ ,  $E_I$ , and  $D$  by minimizing  $\mathcal{L}_L$  in (1),  $\mathcal{L}_S$  in (4), and  $\mathcal{L}_A$  in (9).

Type	Method	OCI→M		OMI→C		OCM→I		ICM→O		#param.	FPS
		HTER	AUC	HTER	AUC	HTER	AUC	HTER	AUC		
1-class	IQM-GMM ( <i>ICB 18</i> )	41.27	55.43	41.84	57.03	39.93	68.99	44.84	36.53	-	31
	OC-fPAD ( <i>IJCB 20</i> )	39.51	60.65	32.64	74.86	38.25	73.01	39.62	69.71	145.03M	210
	OC-LCFAS ( <i>Access 20</i> )	39.10	62.03	43.79	58.43	40.95	49.42	43.32	41.04	8.86M	388
	AAE ( <i>CCBR 21</i> )	42.39	57.29	46.44	46.53	45.07	23.28	43.08	47.93	2.42M	816
	OC-SCMNet ( <i>CVPR 24</i> )	24.05	75.53	28.02	76.92	21.36	87.29	34.37	69.87	5.92M	373
	SLIP (Ours)	<b>18.81</b>	<b>85.55</b>	<b>23.89</b>	<b>82.73</b>	<b>15.71</b>	<b>89.38</b>	<b>29.15</b>	<b>77.14</b>	171.72M	278

Table 3: Cross-domain testing on seen attack protocols.

**Testing** Given an unknown image  $\mathbf{x}_u$  during the inference stage, we adopt the trained SLIP to obtain the estimated SCM of  $\mathbf{x}_u$  and then measure the detection score  $d(\mathbf{x}_u)$  by

$$d(\mathbf{x}_u) = \frac{\sum_{c=1}^C \sum_{h=1}^H \sum_{w=1}^W |D(E_I(\mathbf{x}_u))|}{C \cdot H \cdot W}, \quad (11)$$

where  $C$ ,  $H$ , and  $W$  refer to the channel number, height, and width of the estimated spoof cue map, respectively. To obtain the threshold of binary classification, we refer to previous FAS methods (Wang et al. 2022; Yu et al. 2020b) and use the Youden Index Calculation (Youden 1950).

## Experiment

### Experiment Setting

**Datasets and Evaluation Metrics** We conduct extensive experiments on the following face anti-spoofing databases: (a) **OULU-NPU** (Boulkenafet et al. 2017) (denoted by **O**), (b) **CASIA-MFSD** (Zhang et al. 2012) (denoted by **C**), (c) **MSU-MFSD** (Wen, Han, and Jain 2015) (denoted by **M**), (d) **Idiap Replay-Attack** (Chingovska, Anjos, and Marcel 2012) (denoted by **I**), (e) **3DMAD** (Erdogmus and Marcel 2014) (denoted by **D**), (f) **HKBU-MARs** (Liu et al. 2016b) (denoted by **H**), (g) **CASIA-SURF** (Yu et al. 2020a) (denoted by **U**), and (h) **PADISI-Face** (Rostami et al. 2021) (denoted by **P**).

To ensure a fair comparison with previous one-class FAS methods, we adopt the same evaluation metrics, including APCER (%) ↓, BPCER (%) ↓, ACER (%) ↓ (Standard 2016), HTER (%) ↓ (Anjos and Marcel 2011), and AUC (%) ↑ to report the performance.

**Network Architecture and Implementation Details** We develop SLIP by using the pretrained contrastive language-image pretraining model (CLIP) (Radford et al. 2021) as the network backbone for the image encoder  $E_I$  and the text encoder  $E_T$ . The latent feature reconstructor  $R$  is built using three convolutional layers, and the SCM estimator  $D$  is constructed using a convolutional block that consists of convolutional layers and GELU activation functions. To train SLIP, we set a constant learning rate of  $1e-5$  with Adam optimizer up to 50 epochs.

### Intra-Domain and Cross-Domain Testings

To evaluate the detection performance as well as generalization capability of the proposed SLIP, we conduct both intra-domain and cross-domain testings and compare our results with recent one-class face anti-spoofing methods, including

IQM-GMM (Nikisins et al. 2018), OC-fPAD (Baweja et al. 2020), OC-LCFAS (Lim et al. 2020), AAE (Huang, Xia, and Shen 2021), and OC-SCMNet (Huang et al. 2024a).

**Intra-Domain Testing** In Table 2, we conduct intra-domain testing on **OULU-NPU** (Boulkenafet et al. 2017), which includes print and replay attacks and utilizes various environments, spoof mediums, and capture devices to design four challenging protocols for evaluating the effectiveness of the anti-spoofing models. By leveraging vision foundation models to learn the material information related to potential spoof attacks from the spoof prompts (e.g., photo and print attacks), we see that the proposed SLIP outperforms all the one-class FAS methods with averagely improved 14.79% in ACER.

**Cross-Domain Testing** In Tables 3, 4, and 5, we conduct cross-domain testing involving both seen and unseen attack types to evaluate the generalization capability of the proposed SLIP.

In Table 3, we conduct leave-one-dataset-out testing on the most commonly used benchmarks and measure the model size and the inference throughput. First, we observe that previous one-class FAS methods tend to learn entangled liveness features, leading to poor performance. In contrast, with the proposed prompt-driven liveness feature disentanglement, our method learns to separate domain-specific and live/spoof features, resulting in disentangled liveness features that achieve promising performance compared to previous one-class FAS methods. Next, we compare using different metrics, including the number of parameters (# param.) and inference FPS, to measure the model size and the inference throughput. In particular, although the Gaussian Mixture Model itself includes only a minimal number of parameters, such as weights, means, and covariances, IQM-GMM proposed calculating time-consuming IQM features to learn the GMM, resulting in slow inference speed. Furthermore, because we adopt the pretrained contrastive language-image pretraining model (CLIP) (Radford et al. 2021) as the network backbone for both the image and text encoders during the training stage, the proposed SLIP has a relatively larger model size. Finally, as the proposed SLIP uses only the image encoder and the spoof cue map decoder during the inference stage, it still maintains a relatively fast inference speed.

In Table 4, we adopt the protocols proposed in (Huang et al. 2024a) to conduct cross-domain testing on the datasets **O**, **C**, **M**, **I**, **D**, **H**, and **U**, and evaluate the results when counteracting the unseen attack types. In particular, the authors in (Huang et al. 2024a) proposed adopting the leave-one-attack-

Type	Method	Unseen 3D mask attacks				Unseen print attacks				Unseen replay attacks			
		OM → DHU		OCMI → DHU		OMD → OCMI		OCMI <sub>DHU</sub> → OCMI		OMD → OCMI		OCMI <sub>DHU</sub> → OCMI	
		HTER	AUC	HTER	AUC	HTER	AUC	HTER	AUC	HTER	AUC	HTER	AUC
2-class	IADG (CVPR 23)	32.89	72.51	36.50	69.49	43.98	56.47	38.56	62.14	43.85	55.75	40.04	64.13
	SAFAS (CVPR 23)	38.22	63.75	34.48	65.33	<b>30.85</b>	<b>75.00</b>	40.09	63.16	39.12	64.99	38.45	66.69
1-class	IQM-GMM (ICB 18)	43.58	46.99	43.82	47.18	40.25	62.02	47.56	41.68	37.61	64.66	48.78	41.85
	OC-fPAD (IJCB 20)	39.35	61.86	42.19	57.47	41.59	61.56	40.41	63.83	48.06	42.45	46.87	41.26
	OC-LCFAS (Access 20)	41.74	56.43	41.64	55.11	46.17	53.45	48.29	50.30	41.32	59.08	46.45	53.71
	AAE (CCBR 21)	42.85	55.97	41.07	55.35	48.50	40.94	42.69	57.21	46.70	53.94	37.60	64.68
	OC-SCMNet (CVPR 24)	24.14	74.81	20.85	85.40	37.44	63.23	28.99	72.21	36.41	63.56	29.61	74.99
	SLIP (Ours)	<b>20.9</b>	<b>86.15</b>	<b>17.66</b>	<b>90.48</b>	31.29	74.85	<b>25.81</b>	<b>78.24</b>	<b>34.54</b>	<b>69.53</b>	<b>27.53</b>	<b>78.2</b>

Table 4: Cross-domain testing on unseen attack protocols.

Type	Method	Funny eye		Paper grasses		Silicone 3Dmask	
		HTER	AUC	HTER	AUC	HTER	AUC
1-class	IQM-GMM (ICB 18)	30.11	68.82	15.01	88.82	22.53	80.33
	OC-fPAD (IJCB 20)	45.23	43.19	43.61	45.72	21.65	77.55
	OC-LCFAS (Access 20)	41.88	55.56	36.23	62.12	29.12	69.68
	AAE (CCBR 21)	45.92	46.25	31.20	72.03	27.00	67.26
	OC-SCMNet (CVPR 24)	28.99	60.46	14.33	92.26	8.40	84.27
	SLIP (Ours)	<b>19.77</b>	<b>81.56</b>	<b>7.02</b>	<b>97.61</b>	<b>3.86</b>	<b>98.32</b>

Table 5: Cross-domain testing on unseen physical adversarial attack protocols.

out strategy to consider 3D mask, print, and replay as the unseen attack type within six protocols. The results in Table 4 show that the proposed SLIP significantly outperforms recent one-class FAS methods and achieves competitive performance compared to state-of-the-art two-class FAS methods, including IADG (Zhou et al. 2023) and SAFAS (Sun et al. 2023). By detecting whether a face is occluded by spoof-attack-related objects, the proposed SLIP effectively learns discriminative one-class representations to effectively distinguish live and spoof faces.

In Table 5, to verify SLIP’s effectiveness against physical adversarial attacks, we conduct an experiment where SLIP was trained on live faces from the training dataset  $\mathbf{P}$  (Rostami et al. 2021) and evaluated on the test live faces and various test attacks, including small funny-eye and paper glasses, and large 3D masks from the same dataset. The results in Table 5 show that SLIP effectively detects large physical adversarial attacks, though detecting smaller ones remains highly challenging.

Loss Terms					OCM → I	
$\mathcal{L}_L$	$\mathcal{L}_S$	$\mathcal{L}_{FD}$	$\mathcal{L}_{FA}$	$\mathcal{L}_A$	HTER	AUC
✓	✓	✓			34.43	67.89
✓	✓	✓			26.64	78.33
✓	✓	✓	✓		23.93	80.14
✓	✓	✓		✓	22.43	82.36
✓	✓	✓	✓	✓	19.50	86.93
✓	✓	✓	✓	✓	<b>15.71</b>	<b>89.38</b>

Table 6: Ablation study on the cross-domain protocol OCM → I, under different loss combinations.

## Ablation Study

**Comparison between Different Loss Terms** In Table 6, we compare using different loss terms to train the proposed SLIP. First, we refer to (Baweja et al. 2020; Huang et al. 2024a) by sampling the Gaussian noises as the pseudo spoof features to replace  $\mathbf{s}$  in Figure 2 and use only  $\mathcal{L}_I + \mathcal{L}_S$  to train  $E_I$  and  $D$  as the baseline. We see that the baseline performance remains poor because Gaussian noises are insufficient to simulate the latent spoof features. Next, by comparing the cases of  $\mathcal{L}_I + \mathcal{L}_S$  vs.  $\mathcal{L}_L + \mathcal{L}_S$ , we observe that incorporating prompt learning to generate spoof prompt features significantly improves the performance to effectively address the challenge proposed by one-class training data. Moreover, by comparing the cases of  $\mathcal{L}_L + \mathcal{L}_S$  vs.  $\mathcal{L}_L + \mathcal{L}_S + \mathcal{L}_{FD}$ , we see that the proposed liveness feature disentanglement indeed facilitates the learning of one-class FAS. However, because  $\mathcal{L}_{FD}$  only constrains  $E_T$  to learn the disentangled liveness features within the branch  $E_T-D$  and does not update  $E_I$ , the performance remains limited. By comparing the cases of  $\mathcal{L}_L + \mathcal{L}_S$  vs.  $\mathcal{L}_L + \mathcal{L}_S + \mathcal{L}_{FA}$ , we see that the model benefits from learning out-of-distribution (OOD) features to learn discriminative liveness features by aggregating live features and distancing itself from spoof features. However, due to the missing of  $\mathcal{L}_{FD}$  to learn disentangled liveness features, the performance remains unsatisfactory when tested in a new application domain. Furthermore, when including complete  $\mathcal{L}_{FA} + \mathcal{L}_{FD}$ , we show that SLIP indeed focuses on learning the disentangled liveness features to reduce the live/spoof-irrelative domain information to improve the performance over the case  $\mathcal{L}_L + \mathcal{L}_S$ . Finally, when further including  $\mathcal{L}_A$ , we see that the performance improvement validates the effectiveness of the augmented spoof-like image features to diversify latent spoof features to achieve the best performance.

**Comparison between Different Spoof Prompts** In Table 8, we compare using spoof prompts composed of different placeholders to train the proposed SLIP under the same loss  $\mathcal{L}_L + \mathcal{L}_S$ . Table 7 shows spoof prompt examples composed of different placeholders. First, we adopt the spoof prompts containing only “spoof adjective (SA)” and the all-ones matrix  $\mathbf{1}$  as the corresponding SCMs to train SLIP as the baseline. Because the pretrained CLIP models remain unknown about the abstract concept of “spoof” from the spoof prompts, the performance of the baseline is rather poor. Next,

Placeholder			Spoof Prompts
SA	PA	OO	
✓			This is a {spoof / fake / ...} face.
✓	✓	✓	This is a {spoof / fake / ...} face modified by covering the {upper / right / ...} part with a {mask / photo / ...}.

Table 7: Examples of spoof prompts composed of different placeholders, including spoof adjective (SA), position adjective (PA), and occluding object (OO).

Spoof Prompts			OCM $\rightarrow$ I	
SA	PA	OO	HTER	AUC
✓			30.92	71.12
✓	✓	✓	<b>26.64</b>	<b>78.33</b>

Table 8: Ablation study on the cross-domain protocol OCM  $\rightarrow$  I using spoof prompts composed of different placeholders, including spoof adjective (SA), position adjective (PA), and occluding object (OO).

since live faces should not be covered by any attack-related objects (*e.g.*, paper, screen, or mask), to simulate potential attacks, we further include “position adjective (PA)” and “occluding object (OO)” to construct the spoof prompts as well as the corresponding  $\tilde{\mathbf{m}}$  in Figure 3 to train SLIP. Note that, the vision foundation models (*e.g.*, pretrained CLIP) already possess learned representations of real-world objects. By simulating faces covered by different objects with material characteristics similar to spoof attacks (*e.g.*, paper and print attacks), SLIP learns the concepts of “live” and “spoof” from the designed prompts to effectively distinguish between live and spoof images

## Visualization

**t-SNE Visualization** In Figure 6, we apply *t*-SNE (Van der Maaten and Hinton 2008) to visualize the latent liveness features extracted from  $E_I$  and  $E_T$  with (a) pretrained weights of CLIP model and (b) re-trained weights. In Figure 6(a), because  $E_I$  and  $E_T$  using the initial weights from the CLIP model recognize the concept of “face” but lack an understanding of the concepts of “live” and “spoof”, we see that the latent features extracted from live, spoof, and content prompts are mixed together. In contrast, the visualization results in Figure 6(b) show that the different liveness features cluster well individually after training. In addition, we see that the content features (black crosses) benefit from the proposed disentanglement feature learning, as they are distant from other liveness features in the latent feature space. Finally, observe that some augmented spoof-like image features (yellow dots) diversify latent spoof features to support the FAS models in generalizing to unseen attacks.

**Activation Visualization** In Figure 7, we use Grad-CAM (Selvaraju et al. 2017) to visualize the activation maps of spoof prompts with different position adjectives. To visualize the activation maps, we use  $E_I$  and  $E_T$  to extract  $\mathbf{z}$ ,  $\mathbf{l}$ ,  $\mathbf{s}$  from live images and different live/spoof prompts as the inputs for




Figure 6: *t*-SNE visualizations on the protocol C  $\rightarrow$  I, using (a) pretrained weights of CLIP model and (b) re-trained weights. DOTS: images (green: live, red: spoof) CROSS: prompts (green: live prompt, red: spoof prompt, gray: content prompt, black: hybrid prompt).

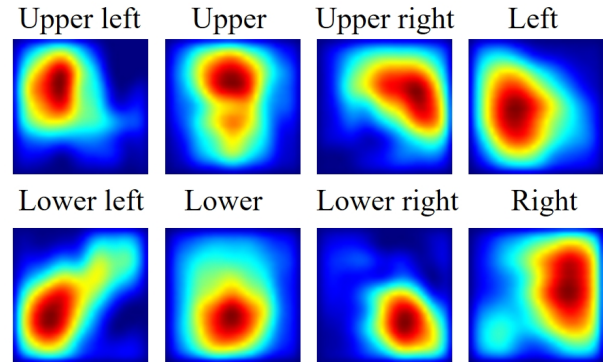


Figure 7: Examples of activation maps, using different position adjectives within spoof prompts.

an additional classifier, which is constrained by cross entropy loss. Observe from Figure 7 that the proposed language-guided spoof cue map estimation indeed adapt  $E_T$  to capture corresponding spoof responses on specified regions from spoof prompts.

## Conclusion

We have introduced a novel Spoof-aware one-class face anti-spoofing with Language Image Pretraining (SLIP) to address the one-class FAS problem. In SLIP, to address the absence of spoof training data, we first propose an effective language-guided spoof cue map estimation by simulating whether the underlying faces are covered by attack-related objects and generating corresponding nonzero spoof cue maps. Next, we first propose a novel prompt-driven liveness feature disentanglement to learn the disentangled liveness features to alleviate live/spoof-irrelative domain variations. Finally, we propose augmenting the spoof-like image features to diversify latent spoof features to facilitate the learning of one-class face anti-spoofing. Extensive experiments demonstrate that SLIP outperforms previous one-class FAS methods.

## Acknowledgments

This work was supported in part by the NSTC grants 111-2221-E-007-109-MY3, 112-2221-E-007-082-MY3, and 113-2634-F-007-002 of Taiwan.

## References

- Anjos, A.; and Marcel, S. 2011. Counter-measures to photo attacks in face recognition: a public database and a baseline. In *2011 international joint conference on Biometrics (IJCB)*, 1–7. IEEE.
- Baweja, Y.; Oza, P.; Perera, P.; and Patel, V. M. 2020. Anomaly detection-based unknown face presentation attack detection. In *2020 IEEE International Joint Conference on Biometrics (IJCB)*, 1–9. IEEE.
- Boulkenafet, Z.; Komulainen, J.; Li, L.; Feng, X.; and Hadid, A. 2017. OULU-NPU: A mobile face presentation attack database with real-world variations. In *2017 12th IEEE international conference on automatic face & gesture recognition (FG 2017)*, 612–618. IEEE.
- Chen, B.; Yang, W.; Li, H.; Wang, S.; and Kwong, S. 2021. Camera invariant feature learning for generalized face anti-spoofing. *IEEE Transactions on Information Forensics and Security*, 16: 2477–2492.
- Chingovska, I.; Anjos, A.; and Marcel, S. 2012. 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)*, 1–7. IEEE.
- Erdogmus, N.; and Marcel, S. 2014. Spoofing face recognition with 3D masks. *IEEE transactions on information forensics and security*, 9(7): 1084–1097.
- Fang, H.; Liu, A.; Jiang, N.; Lu, Q.; Zhao, G.; and Wan, J. 2024. VL-FAS: Domain Generalization via Vision-Language Model For Face Anti-Spoofing. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4770–4774. IEEE.
- Huang, P.-K.; Chang, C.-L.; Ni, H.-Y.; and Hsu, C.-T. 2022a. Learning to augment face presentation attack dataset via disentangled feature learning from limited spoof data. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.
- Huang, P.-K.; Chiang, C.-H.; Chen, T.-H.; Chong, J.-X.; Liu, T.-L.; and Hsu, C.-T. 2024a. One-Class Face Anti-spoofing via Spoof Cue Map-Guided Feature Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Huang, P.-K.; Chiang, C.-H.; Chong, J.-X.; Chen, T.-H.; Ni, H.-Y.; and Hsu, C.-T. 2023a. LDCformer: Incorporating Learnable Descriptive Convolution to Vision Transformer for Face Anti-Spoofing. In *2023 IEEE International Conference on Image Processing (ICIP)*. IEEE.
- Huang, P.-K.; Chin, M.-C.; and Hsu, C.-T. 2021. Face anti-spoofing via robust auxiliary estimation and discriminative feature learning. In *Asian Conference on Pattern Recognition*, 443–458. Springer.
- Huang, P.-K.; Chong, J.-X.; Hsu, M.-T.; Hsu, F.-Y.; Chiang, C.-H.; Chen, T.-H.; Hsu, C.-T.; et al. 2024b. A Survey on Deep Learning-based Face Anti-Spoofing. *APSIPA Transactions on Signal and Information Processing*, 13(1).
- Huang, P.-K.; Chong, J.-X.; Ni, H.-Y.; Chen, T.-H.; and Hsu, C.-T. 2023b. Towards diverse liveness feature representation and domain expansion for cross-domain face anti-spoofing. In *2023 IEEE International Conference on Multimedia and Expo (ICME)*, 1199–1204. IEEE.
- Huang, P.-K.; Lu, C.-Y.; Chang, S.-J.; Chong, J.-X.; and Hsu, C.-T. 2023c. Test-Time Adaptation for Robust Face Anti-Spoofing. In *BMVC*.
- Huang, P.-K.; Ni, H.-Y.; Ni, Y.-Q.; and Hsu, C.-T. 2022b. Learnable Descriptive Convolutional Network for Face Anti-Spoofing. In *BMVC*.
- Huang, X.; Xia, J.; and Shen, L. 2021. One-class face anti-spoofing based on attention auto-encoder. In *Biometric Recognition: 15th Chinese Conference, CCBR 2021, Shanghai, China, September 10–12, 2021, Proceedings 15*, 365–373. Springer.
- Lim, S.; Gwak, Y.; Kim, W.; Roh, J.-H.; and Cho, S. 2020. One-class learning method based on live correlation loss for face anti-spoofing. *IEEE Access*, 8: 201635–201648.
- Liu, A.; Xue, S.; Gan, J.; Wan, J.; Liang, Y.; Deng, J.; Escalera, S.; and Lei, Z. 2024. CFPL-FAS: Class Free Prompt Learning for Generalizable Face Anti-spoofing. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Liu, S.; Yuen, P. C.; Zhang, S.; and Zhao, G. 2016a. 3D Mask Face Anti-spoofing with Remote Photoplethysmography. In *ECCV*, 85–100.
- Liu, S.; Yuen, P. C.; Zhang, S.; and Zhao, G. 2016b. 3D mask face anti-spoofing with remote photoplethysmography. In *European Conference on Computer Vision*, 85–100. Springer.
- Liu, S.-Q.; Lan, X.; and Yuen, P. C. 2018. Remote photoplethysmography correspondence feature for 3d mask face presentation attack detection. In *ECCV*, 558–573.
- Liu, S.-Q.; Lan, X.; and Yuen, P. C. 2021. Multi-Channel Remote Photoplethysmography Correspondence Feature for 3D Mask Face Presentation Attack Detection. *IEEE Transactions on Information Forensics and Security*, 16: 2683–2696.
- Liu, S.-Q.; Lan, X.; and Yuen, P. C. 2022. Learning temporal similarity of remote photoplethysmography for fast 3D mask face presentation attack detection. *IEEE Transactions on Information Forensics and Security*, 17: 3195–3210.
- Liu, S.-Q.; Wang, Q.; and Yuen, P. C. 2024. Bottom-Up Domain Prompt Tuning for Generalized Face Anti-Spoofing. In *European conference on computer vision*. Springer.
- Nikisins, O.; Mohammadi, A.; Anjos, A.; and Marcel, S. 2018. On effectiveness of anomaly detection approaches against unseen presentation attacks in face anti-spoofing. In *2018 International Conference on Biometrics (ICB)*, 75–81. IEEE.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.

- Rostami, M.; Spinoulas, L.; Hussein, M.; Mathai, J.; and Abd-Almageed, W. 2021. Detection and continual learning of novel face presentation attacks. In *Proceedings of the IEEE/CVF international conference on computer vision*, 14851–14860.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- Srivatsan, K.; Naseer, M.; and Nandakumar, K. 2023. Flip: Cross-domain face anti-spoofing with language guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19685–19696.
- Standard, I. 2016. Information Technology—Biometric Presentation Attack Detection—Part 1: Framework. *ISO: Geneva, Switzerland*.
- Sun, Y.; Liu, Y.; Liu, X.; Li, Y.; and Chu, W.-S. 2023. Re-thinking Domain Generalization for Face Anti-spoofing: Separability and Alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24563–24574.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Wang, G.; Han, H.; Shan, S.; and Chen, X. 2020a. Cross-domain face presentation attack detection via multi-domain disentangled representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6678–6687.
- Wang, X.; Zhang, K.-Y.; Yao, T.; Zhou, Q.; Ding, S.; Dai, P.; and Ji, R. 2024. TF-FAS: Twofold-Element Fine-Grained Semantic Guidance for Generalizable Face Anti-Spoofing. In *European Conference on Computer Vision (ECCV)*. Springer.
- Wang, Z.; Wang, Z.; Yu, Z.; Deng, W.; Li, J.; Gao, T.; and Wang, Z. 2022. Domain Generalization via Shuffled Style Assembly for Face Anti-Spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4123–4133.
- Wang, Z.; Yu, Z.; Zhao, C.; Zhu, X.; Qin, Y.; Zhou, Q.; Zhou, F.; and Lei, Z. 2020b. Deep spatial gradient and temporal depth learning for face anti-spoofing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5042–5051.
- Wen, D.; Han, H.; and Jain, A. K. 2015. Face spoof detection with image distortion analysis. *IEEE Transactions on Information Forensics and Security*, 10(4): 746–761.
- Yeh, C.-H.; Hong, C.-Y.; Hsu, Y.-C.; Liu, T.-L.; Chen, Y.; and LeCun, Y. 2022. Decoupled contrastive learning. In *European conference on computer vision*, 668–684. Springer.
- Youden, W. J. 1950. Index for rating diagnostic tests. *Cancer*, 3(1): 32–35.
- Yu, Z.; Cai, R.; Li, Z.; Yang, W.; Shi, J.; and Kot, A. C. 2024. Benchmarking joint face spoofing and forgery detection with visual and physiological cues. *IEEE Transactions on Dependable and Secure Computing*.
- Yu, Z.; Qin, Y.; Zhao, H.; Li, X.; and Zhao, G. 2021. Dual-Cross Central Difference Network for Face Anti-Spoofing. In *IJCAI International Joint Conference on Artificial Intelligence*.
- Yu, Z.; Wan, J.; Qin, Y.; Li, X.; Li, S. Z.; and Zhao, G. 2020a. NAS-FAS: Static-dynamic central difference network search for face anti-spoofing. *IEEE transactions on pattern analysis and machine intelligence*, 43(9): 3005–3023.
- Yu, Z.; Zhao, C.; Wang, Z.; Qin, Y.; Su, Z.; Li, X.; Zhou, F.; and Zhao, G. 2020b. Searching central difference convolutional networks for face anti-spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5295–5305.
- Zhang, Z.; Yan, J.; Liu, S.; Lei, Z.; Yi, D.; and Li, S. Z. 2012. A face antispoofing database with diverse attacks. In *2012 5th IAPR international conference on Biometrics (ICB)*, 26–31. IEEE.
- Zhou, Q.; Zhang, K.-Y.; Yao, T.; Lu, X.; Ding, S.; and Ma, L. 2024. Test-Time Domain Generalization for Face Anti-Spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zhou, Q.; Zhang, K.-Y.; Yao, T.; Lu, X.; Yi, R.; Ding, S.; and Ma, L. 2023. Instance-Aware Domain Generalization for Face Anti-Spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20453–20463.
- Zhou, Q.; Zhang, K.-Y.; Yao, T.; Yi, R.; Ding, S.; and Ma, L. 2022a. Adaptive mixture of experts learning for generalizable face anti-spoofing. In *Proceedings of the 30th ACM International Conference on Multimedia (ACM MM)*, 6009–6018.
- Zhou, Q.; Zhang, K.-Y.; Yao, T.; Yi, R.; Sheng, K.; Ding, S.; and Ma, L. 2022b. Generative domain adaptation for face anti-spoofing. In *European Conference on Computer Vision (ECCV)*, 335–356. Springer.