

Interpretable Face Anti-Spoofing: Enhancing Generalization with Multimodal Large Language Models

Guosheng Zhang^{*1}, Keyao Wang^{*1}, Haixiao Yue^{*1}, Ajian Liu^{†2}, Gang Zhang¹,
Kun Yao¹, Errui Ding¹, Jingdong Wang¹

¹Department of Computer Vision Technology (VIS), Baidu Inc

²CBSR&MAIS, Institute of Automation, Chinese Academy of Sciences (CASIA)

{zhangguosheng, wangkeyao, yuehaixiao, zhanggang03, yaokun, dingerrui}@baidu.com, ajian.liu@ia.ac.cn, wangjingdong@outlook.com

Abstract

Face Anti-Spoofing (FAS) is essential for ensuring the security and reliability of facial recognition systems. Most existing FAS methods are formulated as binary classification tasks, providing confidence scores without interpretation. They exhibit limited generalization in out-of-domain scenarios, such as new environments or unseen spoofing types. In this work, we introduce a multimodal large language model (MLLM) framework for FAS, termed Interpretable Face Anti-Spoofing (I-FAS), which transforms the FAS task into an interpretable visual question answering (VQA) paradigm. Specifically, we propose a Spoof-aware Captioning and Filtering (SCF) strategy to generate high-quality captions for FAS images, enriching the model’s supervision with natural language interpretations. To mitigate the impact of noisy captions during training, we develop a Lopsided Language Model (L-LM) loss function that separates loss calculations for judgment and interpretation, prioritizing the optimization of the former. Furthermore, to enhance the model’s perception of global visual features, we design a Globally Aware Connector (GAC) to align multi-level visual representations with the language model. Extensive experiments on standard and newly devised One to Eleven cross-domain benchmarks, comprising 12 public datasets, demonstrate that our method significantly outperforms state-of-the-art methods.

Introduction

Facial recognition technology has become increasingly sophisticated and is widely utilized for its inherent convenience and contactless operation. These systems are effectively employed in various applications, particularly in on-line payment and identity verification. However, they remain vulnerable to environmental fluctuations and diversified spoofing types, such as printed images, video replays (Boulkenafet et al. 2017), and even high-fidelity 3D masks (Liu et al. 2022a). Consequently, face anti-spoofing is developed to enhance the security and effectiveness of facial recognition systems in a variety of applications. Previ-

^{*}These authors contributed equally.

[†]Corresponding author

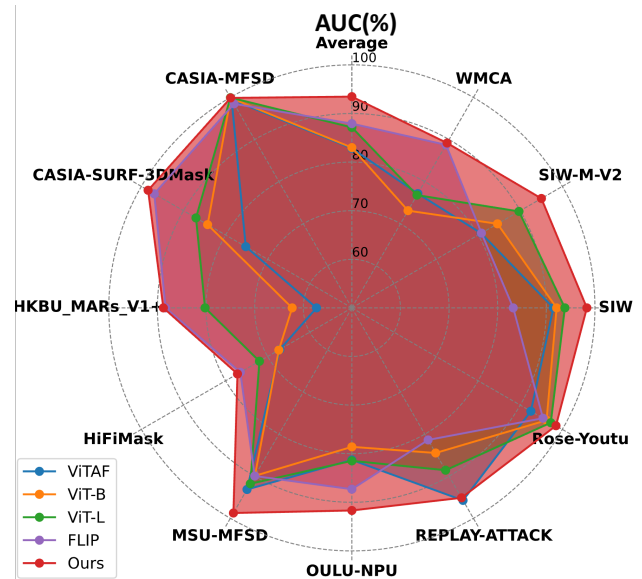


Figure 1: Comparison of performance under a challenging One to Eleven benchmark, where training is restricted to a single source domain (CelebA-Spoof) while testing across 11 target domains. The graph illustrates the notable superiority of our method (red point) compared to existing methods under the condition of a limited source domain.

ous FAS methods (Liu, Jourabloo, and Liu 2018; Yu et al. 2020b) have demonstrated significant achievements, particularly in intra-domain scenarios, but they generally exhibit poor generalization when applied to cross-domain contexts, encountering novel environments, or spoof instruments.

One prevalent approach to improving cross-domain generalization is Domain Adaptation (DA) (Wang et al. 2021; Liu et al. 2022c; Zhou et al. 2022; Yue et al. 2023; Liu et al. 2024d), which aims to reduce the discrepancy between the source and target domains. However, these methods often require access to target domain data during training, which is not always feasible. Another alternative approach is Domain Generalization (DG) (Cai et al. 2022; Zhou et al. 2023;

Cai et al. 2024; Liu 2024; Le and Woo 2024; Zhou et al. 2024), designed to train models that can generalize robustly to the target domain by accessing multiple source domains. DG-based methods typically learn a theoretically domain-invariant feature representation across multiple source domains, employing techniques such as adversarial training or feature disentanglement. However, the substantial distributional differences between source domains, coupled with the practical challenges of diverse spoofing types, render the reliance on category-level annotations alone insufficient for the model to learn domain-invariant features.

Recent research (Srivatsan, Naseer, and Nandakumar 2023; Liu et al. 2024b,a; Guo et al. 2024a; Shi et al. 2024) has proposed employing a CLIP-like framework for FAS. This approach outperforms the traditional unimodal approach by incorporating a textual modality that provides descriptive textual information for FAS images, thereby aiding in model training and decision-making. Although utilizing semantic content from the text significantly enhances the model’s generalization capabilities, the reliance on manually constructed text based on prior knowledge leads to limited diversity and a scarcity of instructive information.

Rethinking how humans effortlessly and robustly identify spoofs, we observe that they can disregard irrelevant factors, concentrate on key spoofing cues, and derive interpretable judgments through causal reasoning. Inspired by recent breakthroughs in multimodal large language models (MLLMs), which have demonstrated exceptional image-text comprehension capabilities and strong generalization capabilities. We propose a pioneering MLLM framework for FAS, termed Interpretable FAS (I-FAS), which transforms the FAS classification task into an interpretable framework of Visual Question Answering (VQA), providing additional natural language interpretations for judgments. However, current FAS datasets typically rely on category-level annotations, lacking granular details such as comprehensive image descriptions, especially those that highlight spoofing clues. To bridge this gap, we design the Spoof-aware Captioning and Filtering (SCF) strategy. This strategy integrates two distinct captioners: a general captioner that generates captions for real samples, and a spoof-aware captioner that specializes in providing captions with spoof-specific preferences, such as the type, medium, and form. To mitigate the impact of noisy captions during training, we develop a Lopsided Language Model (L-LM) loss function that differentiates between the judgment and the interpretation components of the answer, allowing for separate loss calculations. By increasing the loss weight attributed to the judgment, we substantially accelerate and enhance the stability of model convergence. Furthermore, to enhance the FAS model’s perception of global visual features, particularly focusing on low-level visual features like moiré patterns from screens and blur associated with paper, we introduce the Globally Aware Connector (GAC). The GAC integrates multi-level global representation, providing the language model with a comprehensive understanding from the global to the local perspectives. Finally, to better demonstrate the generalizability of our approach, we have established a more extensive and challenging benchmark, One to Eleven, including

12 public FAS datasets. Our method elucidates the rationale behind its decisions, thereby enhancing interpretability and robustness, and improving cross-domain generalization performance.

- We reformulate the FAS classification task into an Interpretable VAQ paradigm, termed I-FAS, which leverages a Globally Aware Connector (GAC) to capture multi-level spoof cues, thereby enhancing robustness and cross-domain generalization.
- We introduce Spoof-aware Captioning and Filtering (SCF) to provide more comprehensive annotation data for FAS, enabling the model to explain its decision-making process, thereby improving interpretability.
- We developed a Lopsided LM (L-LM) loss that separates the answer into judgment and interpretation for separate loss computation, mitigating the impact of noisy captions during training.
- Extensive experiments on both standard and newly designed (One to Eleven) cross-domain benchmarks demonstrate that our method achieves a significant improvement over state-of-the-art methods.

Related Work

Face Anti-Spoofing

FAS aims to determine whether an image captures a genuine or a deceptive presentation attack. Early research was rooted in traditional handcrafted features and evolved into more advanced deep learning techniques (Wang et al. 2022a; Zhang et al. 2020a; Wang et al. 2024a). (Yu et al. 2020b) have focused on designing specialized architectures for FAS, including the recently superior effective transformer structures highlighted by (Huang et al. 2022). Furthermore, advancements have been facilitated through auxiliary supervision signals, such as depth maps (Liu, Jourabloo, and Liu 2018) and reflection maps (Zhang et al. 2021). Despite their success in intra-dataset scenarios, these methods often falter in cross-domain settings. To address this issue, recent methods have employed DA-based techniques (Wang et al. 2021; Liu et al. 2022c; Zhou et al. 2022; Yue et al. 2023; Liu et al. 2024d) to reduce inter-domain distribution discrepancies by introducing unlabeled target domain data. Concurrently, DG-based approaches (Zhou et al. 2023; Cai et al. 2024; Liu 2024; Le and Woo 2024; Zhou et al. 2024) aim to learn domain-invariant features across multiple source domains via adversarial learning (Shao et al. 2019; Jia et al. 2020), meta-learning (Cai et al. 2022; Chen et al. 2021). Additionally, incremental learning (IL) methods (Hu et al. 2024; Guo et al. 2022; Cai et al. 2023; Wang et al. 2024b) are considered to tackle the catastrophic forgetting problem in the context of domain discontinuity. However, these methods concentrate on extracting generalized liveness-specific features, relying exclusively on binary labels.

Vision-Language Models

Recently, multimodal vision-language models have seen remarkable advancements, delivering promising performance across a spectrum of tasks. Pioneering work (Radford et al.

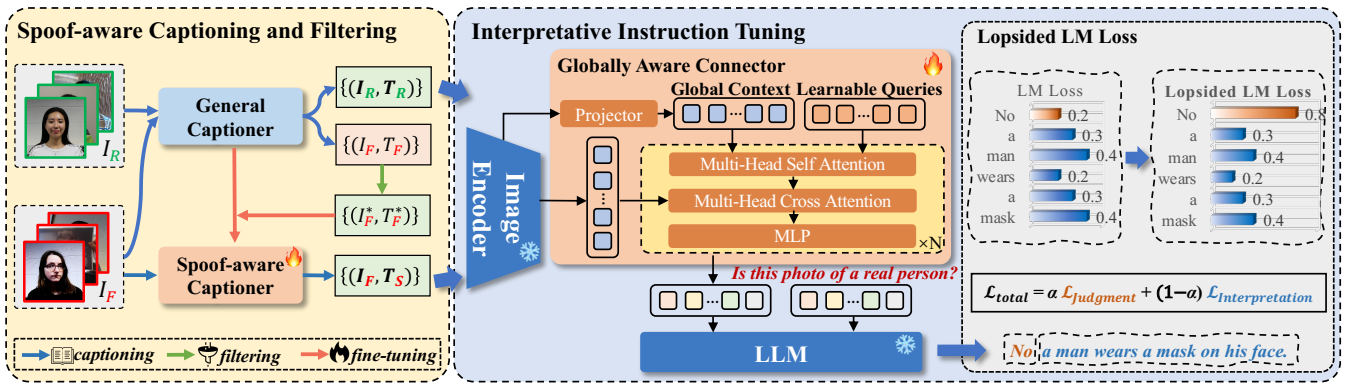


Figure 2: Overview of the proposed Interpretable Face Anti-Spoofing (I-FAS) framework. section illustrates the process of our proposed Spoof-aware Captioning and Filtering (SCF) strategy. The central section details the model architecture, which includes a frozen visual encoder, a pre-trained language model (LLM), and the Globally Aware Connector (GAC). The rightmost section presents a schematic representation of the Lopsided Language Model (L-LM) loss.

2021; Huang et al. 2024; Jiang et al. 2024) focused on vision-language pre-training, which aimed to cultivate foundational models capable of enhanced performance on diverse vision-language tasks. Recent methods (Liu et al. 2024c; Li et al. 2023) have harnessed knowledge from frozen LLMs (Zhang et al. 2022) for vision-to-language generation tasks in an instruction-tuning manner, commonly referred to as multimodal large language models (MLLMs). The pivotal challenge is the design of a cross-modal connector. This connector must align visual features from the pre-trained visual backbone to the word embedding space of the language model. Vision-Language models (VLM) have catalyzed innovative thinking and progress in numerous disciplines (Yuan et al. 2024; Zanella et al. 2024; Guo et al. 2023, 2024b). In terms of FAS, (Srivatsan, Naseer, and Nandakumar 2023) have adapted the multimodal pre-trained CLIP model, grounding visual representations with natural language supervision to enhance generalizability. (Liu et al. 2024b) have designed a CLIP-like model to expand the semantic space through textual prompt learning, thereby fine-tuning visual features for improved generalization. Instead of adopting inherent language templates like (Zhang et al. 2025), our approach leverages interpretable and spoof-aware text to furnish the model with valuable supervisory signals, thereby enriching its learning process and decision-making capabilities.

Methodology

Spoof-aware Captioning and Filtering

Off-the-shelf captioners, trained on a broad range of generic scene images, often focus on general attributes like expressions, clothing, and environment when generating descriptions for FAS images, as shown in Figure 3. They exhibit significant limitations in identifying specific spoof-related cues, such as the type, medium, and form. To address this, we develop a Spoof-aware Captioning and Filtering (SCF) strategy, which aims to generate interpretable descriptions elucidating the rationale for classifying images as either genuine or a presentation attack. As illustrated in Algorithm 1,

we begin by aggregating 12 public FAS datasets into a cohesive dataset $\mathcal{D} = \{(I^i, Y^i)\}_{i=1}^N$, where I^i represents the images categorized as either real (I_R) or fake (I_F) and Y^i denotes category label. Concurrently, we establish a spoof-aware keyword dictionary \mathcal{K} categorized by spoof types, including print, replay, mask, and mannequin attacks. We initiate the process by generating captions for all fake images I_F using an off-the-shelf general captioner C_G (Li et al. 2023), resulting in textual descriptions T_F . We then filter samples, retaining only those samples whose captions contain keywords from \mathcal{K} that correspond to their spoof type, culminating in the formation of a spoof-aware dataset \mathcal{D}_S . For example, if a sample’s caption contains the keyword “screen” and is annotated as a screen attack, we can heuristically infer that the image likely contains discernible cues of the screen attack. Further detailed analysis can be found in the Appendix.

To ensure all fake samples are endowed with spoof-aware captions, we finetune the general captioner C_G using our curated spoof-aware dataset \mathcal{D}_S , resulting in the spoof-aware captioner C_S , which we hypothesize is capable of inherently recognizing and describing spoof-related cues, as shown in Figure 3. For real samples, we again employ C_G to generate captions for all real samples I_R . Finally, by incorporating differential captions corresponding to the original category labels, we achieve a more comprehensive dataset that not only categorizes but also elucidates the rationale underpinning each judgment.

Interpretable Instruction Tuning

Revisiting MLLMs: MLLMs are designed to address sophisticated tasks by generating responses using multimodal inputs, including visual and textual data. The architecture of MLLMs comprises three principal components: 1) the pre-trained visual encoders E_V : This component converts the input image $I \in \mathbb{R}^{H \times W \times 3}$ into a set of visual features $X_V \in \mathbb{R}^{N \times D}$, where N and D denotes the count and dimension of visual features respectively. 2) Vision-Language Connector $P_{V \rightarrow T}$: This element is tasked with aligning visual features to textual token space T of LLMs. 3) LLMs

Algorithm 1: Spoof-aware Captioning and Filtering

Input: Dataset $\mathcal{D} = \{(I^i, Y^i)\}_{i=1}^N$, where $I^i \in \{I_R, I_F\}$, $Y^i \in \{Y_R, Y_F\}$, $F = \{print, replay, mask, mannequin\}$
Keywords: $\mathcal{K} = \{“paper” : Y_{print}, “screen” : Y_{replay}, \dots\}$
General captioner: C_G
Output: Dataset $\mathcal{D}_{cap} = \{(I^i, Y^i, T^i)\}_{i=1}^N$, where $T^i \in \{T_R, T_S\}$

- 1: Captioning $T_F = C_G(I_F)$
- 2: Initialize empty dataset \mathcal{D}_S
- 3: **for** each sample (I_F^i, Y_F^i, T_F^i) **do**
- 4: **for** each keyword k in \mathcal{K} **do**
- 5: **if** k in T_F^i and $\mathcal{K}[k]$ match Y_F^i **then**
- 6: $\mathcal{D}_S \leftarrow \mathcal{D}_S \cup \{(I_F^i, Y_F^i, T_F^i)\}$
- 7: **end if**
- 8: **end for**
- 9: **end for**
- 10: Finetune C_G with \mathcal{D}_S then obtain C_S
- 11: Captioning $T_R = C_G(I_R)$ and $T_S = C_S(I_F)$
- 12: $\mathcal{D}_{cap} \leftarrow \{(I_R, Y_R, T_R)\} \cup \{(I_F, Y_F, T_S)\}$
- 13: **return** \mathcal{D}_{cap}



Figure 3: Illustration of some image-caption pair from the spoof sample. The captions T_F and T_S are generated by general captioner C_G and spoof-aware captioner C_S , respectively. The keywords instrumental in identifying spoof cues are distinctly highlighted in red within the captions.

Φ_T : As the cornerstone of MLLMs, these models are capable of auto-regressively generating free-form responses when prompted with visual tokens X_V and textual tokens X_T . Specifically, for a sequence of length L , the probability of generating target answers X_A is computed by:

$$p(X_A | X_V, X_T) = \prod_{i=1}^L p(X_{A,i} | X_V, X_{T,<i}, X_{A,<i}) \quad (1)$$

In the conventional MLLM framework, the effectiveness of the visual encoder and LLM is largely influenced by the pre-training phase. Consequently, the design of the connector is crucial during the instructing tuning stage, as it is responsible for the extraction and transformation of pertinent information from the visual features into the textual domain, effectively navigating through redundant visual information.

Globally Aware Connector: The vision-language connector is designed to transform the visual features X_V into

textual tokens X_T compatible with textual processing. To enhance the model’s perception of multi-level visual features, we introduce the Globally Aware Connector (GAC), which enhances global perception during visual feature extraction by incorporating multi-level global context features. As illustrated in Figure 2, we utilize multi-layer global visual features $G_V = \{g_1, g_2, \dots, g_L\}$, projected by linear layer, as additional input queries $Q_V \in \mathbb{R}^{L \times D}$, where L is the total layers of the image encoder and g_i represents cls token of the i -th layers. These queries engage with the learnable queries $Q_P \in \mathbb{R}^{M \times D}$ through multi-head self-attention layers (MSA), where M denotes the number of queries. Concurrently, the learnable queries interact with local visual features X_V via multi-head cross-attention operations (MCA). This process facilitates a comprehensive integration of global and local visual information and can be expressed as:

$$\begin{aligned}
 Q &= \text{Concat}(Q_P, Q_V) \\
 Q' &= Q + \text{MSA}(\text{LN}(Q)) \\
 Q'' &= Q' + \text{MCA}(\text{LN}(Q'), \text{LN}(X_V)) \\
 X_T &= Q'' + \text{MLP}(\text{LN}(Q''))
 \end{aligned} \quad (2)$$

Research (Jiang et al. 2023) indicates that different layers of visual encoder exhibit distinct biases towards various patterns: shallow layers are adept at capturing detailed low-level information, while deep layers are proficient in semantic comprehension. By this mechanism, the LLM receives visual information conducive to global perception. The GAC provides LLM with a comprehensive visual representation that integrates both the macroscopic context and the microscopic details.

Training with Lopsided LM Loss: In this work, we reformulate the FAS classification task into the VQA paradigm. For each training instance, we construct single-turn VQA data in the form of (I, T_Q, T_A) , where T_Q is a unified question: “Is this photo of a real person?” serving as instruction. The target answers are formatted as follows:

$$T_A = [T_{Judgment}, T_{Interpretation}] \quad (3)$$

where $T_{Judgment} \in \{“Yes”, “No”\}$ indicates the judgment outcome, while $T_{Interpretation} = \text{This is } \langle T_R/T_S \rangle$ comprises descriptive captions generated by the general captioner for real images (T_R) and the spoof-aware captioner for fake images (T_S).

As shown in Figure 2, to mitigate the impact of noisy interpretations, we introduce a Lopsided Language Model (LLM) Loss that separately calculates the loss for $T_{Judgment}$ and $T_{Interpretation}$, enabling the model to prioritize the accuracy of judgments. The training objective is defined as:

$$\mathcal{L}_{total} = \alpha \mathcal{L}_{Judgment} + (1 - \alpha) \mathcal{L}_{Interpretation} \quad (4)$$

where α is a hyperparameter that balances the emphasis between judgment and interpretation loss components. In the inference phase, the model’s prediction is determined by the probability assigned to the word “Yes”, which indicates the likelihood of the sample being classified as real.

Methods	O&C&I to M		O&M&I to C		O&C&M to I		I&C&M to O		Avg.
	HTER(%)	AUC(%)	HTER(%)	AUC(%)	HTER(%)	AUC(%)	HTER(%)	AUC(%)	HTER(%)
FGHV (Liu et al. 2022b)	9.17	96.92	12.47	93.47	16.29	90.11	13.58	93.55	12.88
GDA (Zhou et al. 2022)	9.20	98.00	12.20	93.00	10.00	96.00	14.40	92.60	11.45
PatchNet (Wang et al. 2022a)	7.10	98.46	11.33	94.58	13.40	95.67	11.82	95.07	10.91
SSAN (Wang et al. 2022b)	6.67	98.75	10.00	96.67	8.88	96.79	13.72	93.63	9.82
IADG (Zhou et al. 2023)	5.41	98.19	8.70	96.40	10.62	94.50	8.86	97.14	8.40
UDG-FAS (Liu et al. 2023)	5.95	98.47	9.82	96.76	5.86	98.62	10.97	95.36	8.15
TTDG (Zhou et al. 2024)	4.16	98.48	7.59	98.18	9.62	98.18	10.00	96.15	7.84
SA-FAS (Sun et al. 2023)	5.95	96.55	8.78	95.37	6.58	97.54	10.00	96.23	7.83
DiVT-M (Liao et al. 2023)	2.86	99.14	8.67	96.92	3.71	99.29	13.06	94.04	7.08
GAC-FAS (Le and Woo 2024)	5.00	97.56	8.20	95.16	4.29	98.87	8.60	97.16	6.52
FLIP (Srivatsan et al. 2023)	4.95	98.11	0.54	99.98	4.25	99.07	2.31	99.63	3.01
CFPL (Liu et al. 2024b)	1.43	99.28	2.56	99.10	5.43	98.41	2.50	99.42	2.98
Ours	0.32	99.88	0.04	99.99	3.22	98.48	1.74	99.66	1.33

Table 1: Comparison with the closest and SOTA FAS methods in Protocol 1 on MSU-MFSD (M), CASIA-FASD (C), Replay-Attack (I), and OULU-NPU (O) datasets. Avg indicates the average performance across four experimental scenarios. The scores presented in bold represent the best performance.

Experiments

Experimental Setup

Databases, Protocols, and Evaluation Metrics: We evaluate our method on two protocols. For Protocol 1, Following established practices, we implement the leave-one-domain-out testing approach on several datasets: MSU-MFSD (M) (Wen, Han, and Jain 2015), CASIA-MFSD (C) (Zhang et al. 2012), Idiap Replay Attack (I) (Chingovska, Anjos, and Marcel 2012), and OULU-NPU (O) (Boulkenafet et al. 2017). To assess the robustness of our method in more demanding conditions, we set up Protocol 2 as One to Eleven testing protocol. Employing only CelebA-Spoof (Zhang et al. 2020b) as the source domain, and 11 datasets as target domains for cross-domain testing. This selection include MSU-MFSD (Wen, Han, and Jain 2015), CASIA-MFSD (Zhang et al. 2012), Idiap Replay Attack (Chingovska, Anjos, and Marcel 2012), OULU-NPU (Boulkenafet et al. 2017), SIW (Liu, Jourabloo, and Liu 2018), Rose-Youtu (Li et al. 2018), HKBU-MARs-V1+ (Liu, Lan, and Yuen 2018), WMCA (George et al. 2019), SIW-M-V2 (Guo et al. 2022), CASIA-SURF-3DMask (Yu et al. 2020a) and HiFiMask (Liu et al. 2022a). For both protocols, we utilize the Area Under the Curve (AUC) and the Half Total Error Rate (HTER) as our primary evaluation metrics. The higher AUC and lower HTER signify superior performance.

Implementation Details: We crop the face images and resize them to $224 \times 224 \times 3$ with RGB channels. For the frozen image encoder, we utilize pre-trained vision models: ViT-L/14 from CLIP (Radford et al. 2021). Following (Li et al. 2023), OPT-2.7B (Zhang et al. 2022) is adopted as the pre-trained large language model. We use the AdamW optimizer, with an initial learning rate set to 10^{-5} and a weight decay parameter set to 10^{-2} . We configure our training process with a batch size of 32 and a maximum of 10 epochs for both Protocol 1 and Protocol 2. For Protocol 2, we meticulously reproduce the baseline methods, including FLIP (Srivatsan, Naseer, and Nandakumar 2023) and ViTAF (Huang et al. 2022), using the official code provided. Both ViT-B

and ViT-L are pre-trained CLIP (Radford et al. 2021). To ensure the integrity and reproducibility of our experiments, we report all results as the mean of three independent runs, each with a unique initialization seed.

Comparison Results

Comparison Results in Protocol 1: We evaluate our method using four leave-one-domain-out settings and compare its performance with that of the latest state-of-the-art approaches. As shown in Table 1, our method demonstrates a significant performance advantage over all single-modal methods, outperforming them by a substantial margin. This substantial advantage highlights the efficacy of multimodal learning in enhancing the model’s generalizability. Furthermore, our approach also surpasses recent multimodal methods, such as (Srivatsan, Naseer, and Nandakumar 2023) and (Liu et al. 2024b), as evidenced by the average HTER reductions to 1.33% from 2.98% and 3.01%. This improvement suggests that the incorporation of richer and more insightful textual information facilitates the acquisition of more robust and widely applicable features by the model.

Comparison Results in Protocol 2: To further evaluate the generalizability of our approach, we utilize a single dataset as the source domain and perform cross-domain testing across 11 distinct datasets. As shown in Table 2, our method demonstrates a compelling advantage of more than 7% relative to the current state-of-the-art techniques. This substantial advantage is particularly evident under severely constrained source domain conditions. This scenario reflects real-world challenges that require robustness against emerging attack vectors and shifts in environmental conditions. Significantly, the datasets CASIA-SURF-3DMask, HKBU-MARs-V1+, and SIW-M-V2 encompass attack modalities absent in the source domain (CelebA-Spoof), including sophisticated 3D attacks, makeup alterations, and novel material-based attacks. Under these challenging conditions, our method consistently outperforms conventional single-modal DG models, demonstrating the

(a) Average Over 11 Datasets			(b) CASIA-MFSD			(c) CASIA-SURF-3DMask			(d) HKBU-MARs-V1+		
Methods	HTER(%)	AUC(%)	Methods	HTER(%)	AUC(%)	Methods	HTER(%)	AUC(%)	Methods	HTER(%)	AUC(%)
ViTAF	23.85	82.82	ViTAF	3.11	99.48	ViTAF	32.44	75.20	ViTAF	49.29	57.28
ViT-B	23.48	82.98	ViT-B	0.70	99.86	ViT-B	24.89	84.26	ViT-B	45.08	62.28
ViT-L	21.08	85.61	ViT-L	0.93	99.95	ViT-L	23.54	84.22	ViT-L	33.33	73.88
FLIP	18.73	87.90	FLIP	4.88	98.48	FLIP	8.83	96.93	FLIP	17.25	88.31
Ours	11.30	93.71	Ours	1.11	99.88	Ours	6.18	98.40	Ours	18.64	88.77

(e) HiFiMask			(f) MSU-MFSD			(g) OULU-NPU			(h) REPLAY-ATTACK		
Methods	HTER(%)	AUC(%)	Methods	HTER(%)	AUC(%)	Methods	HTER(%)	AUC(%)	Methods	HTER(%)	AUC(%)
ViTAF	37.30	67.10	ViTAF	12.86	93.14	ViTAF	26.73	81.28	ViTAF	12.38	95.73
ViT-B	37.33	67.35	ViT-B	16.67	89.89	ViT-B	28.53	78.59	ViT-B	24.80	84.47
ViT-L	32.81	72.58	ViT-L	20.87	85.65	ViT-L	29.42	78.07	ViT-L	16.58	92.00
FLIP	28.32	76.50	FLIP	19.37	89.95	FLIP	20.57	87.30	FLIP	25.67	81.37
Ours	28.23	77.17	Ours	5.63	98.73	Ours	14.86	91.68	Ours	9.15	95.12

(i) Rose-Youtu			(j) SIW			(k) SIW-M-V2			(l) WMCA		
Methods	HTER(%)	AUC(%)	Methods	HTER(%)	AUC(%)	Methods	HTER(%)	AUC(%)	Methods	HTER(%)	AUC(%)
ViTAF	69.34	74.22	ViTAF	14.74	92.51	ViTAF	26.72	80.70	ViTAF	29.88	77.14
ViT-B	82.69	63.22	ViT-B	9.13	96.24	ViT-B	22.60	84.59	ViT-B	34.72	73.10
ViT-L	80.47	71.69	ViT-L	9.03	96.56	ViT-L	17.26	90.37	ViT-L	34.39	75.13
FLIP	80.73	73.60	FLIP	11.01	95.40	FLIP	25.95	80.78	FLIP	19.36	88.73
Ours	5.52	98.48	Ours	4.02	98.34	Ours	10.89	95.02	Ours	20.07	89.17

Table 2: Comparison in Protocol 2, illustrating the challenge of training solely on the CelebA-Spoof dataset followed by testing across 11 distinct datasets. We run each experiment 3 times under different seeds and report the average HTER and AUC.

SCF	GAC	L-LM	HTER(%)	AUC(%)
	✓	✓	14.47	90.41
✓		✓	14.84	90.78
✓	✓		12.06	92.81
✓	✓	✓	11.30	93.71

Table 3: Ablation study on the effectiveness of key components within our proposed method, including the Spoof-aware Captioning and Filtering (SCF), Globally Aware Connector (GAC), and Lopsided LM (L-LM) Loss. The results are the average HTER and AUC in Protocol 2.

model’s superior adaptability and resilience.

Ablation Study

Study on Each Component. To systematically assess the impact of individual components within our framework, we perform an ablation study of the SCF, the GAC, and the L-LM Loss. For each component, we conduct sequential removal experiments within the framework and report the average results for HTER and AUC in Protocol 2.

As shown in Table 3, the first row corresponds to the scenario where the SCF strategy was removed, replacing the interpretable target answer with a simple binary judgment (“Yes/No”). This modification resulted in a significant performance degradation, with an increase of 3.17% in HTER and a decrease of 3.30% in AUC. The second row indicates the impact of excluding the GAC module. To maintain experimental integrity and a consistent token number for LLMs, we replaced the multi-level global features

Methods	Text Format	HTER(%)	AUC(%)
FLIP	$T_{Template}$	18.73	87.90
FLIP	$T_{Judgment, \{T_R, T_S\}}$	17.30	89.15
Ours	$T_{Judgment}$	14.47	90.41
Ours	$T_{Judgment, \{T_R, T_F\}}$	14.23	90.22
Ours	$T_{Judgment, \{T_R, T_S\}}$	11.30	93.71

Table 4: Ablation study on the effects of various text formats as supervision. The results are the average HTER and AUC in Protocol 2.

with learnable queries. This alteration also culminated in a marked performance drop of +3.54% (HTER) and -2.93% (AUC), respectively. The third row presents the results of employing a standard LM Loss in place of our proposed Lopsided LM Loss. This change likewise resulted in a minor yet discernible reduction in the model’s generalization capability, as indicated by a slight performance drop.

Study on Spoof-aware Interpretation. To highlight the advantages of interpretable annotations, we apply them to the existing multimodal FAS method FLIP (Srivatsan, Naseer, and Nandakumar 2023). Specifically, we only replace the original manual text template with spoof-aware interpretations. The Table 4 show a significant improvement with a 1.43% decrease in HTER and a 1.25% increase in AUC. To isolate the impact of caption diversity from the interpretability of spoof clues, we further replace the spoof-aware captions T_S with those (T_F) generated by a general captioner devoid of spoof-specific preferences. Comparing

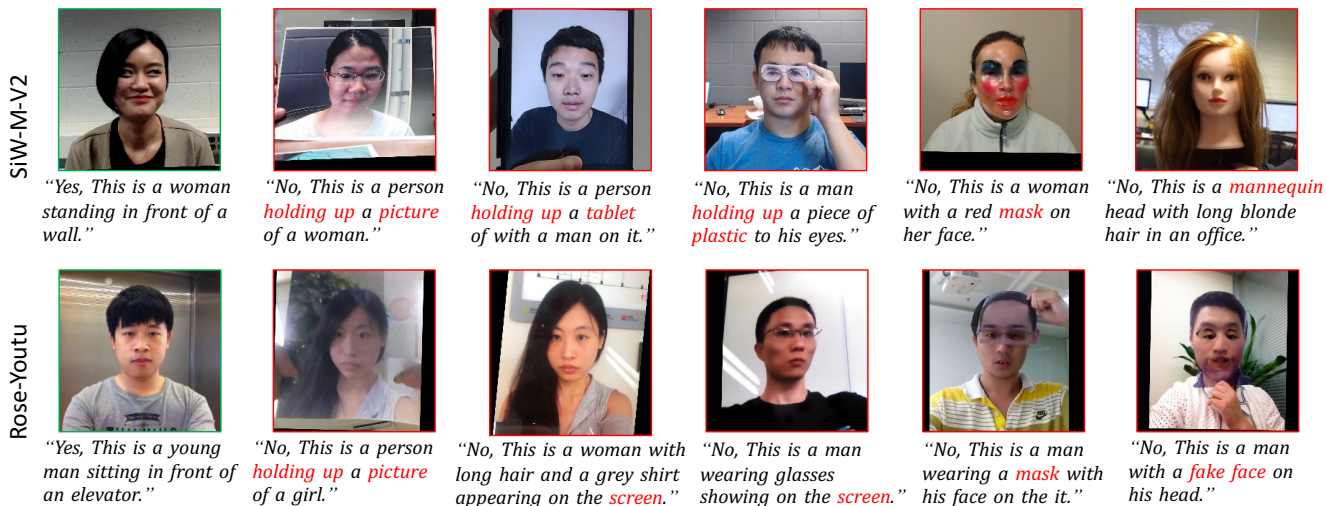


Figure 4: The output responses of I-FAS on some images from SIW-M-V2 and Rose-Youtu dataset under the unified question: “Is this photo of a real person?”. The green box represents the real person, and the red box represents the spoofer sample.




Figure 5: Left: Ablation analysis of hyperparameters α of lopsided LM loss in Protocol 2. Right: Visualization of loss convergence behavior with and without lopsided LM loss.

the results from the second to last row and the third to last row in Table 4, it’s clear that diverse captions alone do not enhance FAS judgments. In contrast, our spoof-aware interpretations are demonstrated to confer substantial improvements in the model’s generalization capabilities.

Study on Lopsided LM Loss. Figure 5 (Left) presents an ablation analysis of the hyperparameter α in lopsided LM loss, which is used to balance the emphasis between judgment and interpretation loss components. the average and variance of the AUC and HTER in Protocol 2 across a range of α values. Our findings indicate that extreme values of α lead to performance degradation. We suggest that an overly large α may diminish the generalization benefits conferred by interpretation components, as it disproportionately emphasizes the judgment aspect. Conversely, placing excessive weight on interpretation may introduce noise, thereby hindering the model’s ability to learn effectively. As shown in Figure 5 (Right), not appropriately increasing the importance of judgment through the lopsided LM loss can result in slow and unstable model convergence during training.

Visualization and Analysis

Figure 3 displays a selection of image-text pairs from the CelebA-Spoof dataset, crafted by our SCF strategy. Captions T_F generated by the general captioner C_G offer generic descriptions, capturing general attributes without bias toward specific features. In contrast, captions T_S crafted by the spoof-aware captioner C_S reveal the underlying spoofing tactics by emphasizing keywords such as “paper”, “tablet”, and “screen”. Such annotations provide invaluable supervisory signals that are instrumental in the learning process of FAS tasks, enabling the model to discern subtle yet critical cues indicative of spoofing attempts.

In Protocol 2, after training solely on the CelebA-Spoof dataset, our model is tested on multiple unseen datasets. Figure 4 illustrates the model’s output responses post-interpretative instruction tuning. Our method excels not only in rendering accurate judgments but also in furnishing rationales for these decisions, especially for samples with attack behaviors. The dual capability of accurate classification and interpretability strengthens the robustness and generalizability of our approach, which is pivotal for real-world applications where understanding the rationale for the decisions is as crucial as the decisions themselves.

Conclusion

In this work, we introduce a novel perspective for FAS by transforming it into an interpretable VQA task. Incorporating SCF strategies allows the model to generate interpretive captions, significantly enhancing cross-domain generalization capabilities. Our proposed Lopsided Language Model Loss optimizes the training process. Furthermore, the GAC enhances the model’s perception of multi-level global context features, providing substantial performance improvements for liveness-specific feature learning. Experiments on Protocols 1 and 2 demonstrate the exceptional performance of our method across multiple public FAS datasets.

References

- 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.
- Cai, R.; Cui, Y.; Li, Z.; Yu, Z.; Li, H.; Hu, Y.; and Kot, A. 2023. Rehearsal-free domain continual face anti-spoofing: Generalize more and forget less. In *CVPR*, 8037–8048.
- Cai, R.; Li, Z.; Wan, R.; Li, H.; Hu, Y.; and Kot, A. C. 2022. Learning meta pattern for face anti-spoofing. *TIFS*, 17: 1201–1213.
- Cai, R.; Soh, C.; Yu, Z.; Li, H.; Yang, W.; and Kot, A. C. 2024. Towards Data-Centric Face Anti-spoofing: Improving Cross-Domain Generalization via Physics-Based Data Synthesis. *International Journal of Computer Vision*, 1–22.
- Chen, Z.; Yao, T.; Sheng, K.; Ding, S.; Tai, Y.; Li, J.; Huang, F.; and Jin, X. 2021. Generalizable representation learning for mixture domain face anti-spoofing. In *AAAI*, volume 35, 1132–1139.
- Chingovska, I.; Anjos, A.; and Marcel, S. 2012. On the effectiveness of local binary patterns in face anti-spoofing. In *BIOSIG*, 1–7. IEEE.
- George, A.; Mostaani, Z.; Geissenbuhler, D.; Nikisins, O.; Anjos, A.; and Marcel, S. 2019. Biometric face presentation attack detection with multi-channel convolutional neural network. *TIFS*, 15: 42–55.
- Guo, J.; Liu, H.; Luo, Y.; Hu, X.; Zou, H.; Zhang, Y.; Liu, H.; and Zhao, B. 2024a. Style-conditional Prompt Token Learning for Generalizable Face Anti-spoofing. In *ACM MM*, 994–1003.
- Guo, X.; Liu, X.; Masi, I.; and Liu, X. 2024b. Language-guided Hierarchical Fine-grained Image Forgery Detection and Localization. *International Journal of Computer Vision*.
- Guo, X.; Liu, X.; Ren, Z.; Grosz, S.; Masi, I.; and Liu, X. 2023. Hierarchical fine-grained image forgery detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3155–3165.
- Guo, X.; Liu, Y.; Jain, A.; and Liu, X. 2022. Multi-domain learning for updating face anti-spoofing models. In *European Conference on Computer Vision*, 230–249. Springer.
- Hu, C.; Zhang, K.-Y.; Yao, T.; Liu, S.; Ding, S.; Tan, X.; and Ma, L. 2024. Domain-Hallucinated Updating for Multi-Domain Face Anti-spoofing. In *AAAI*, volume 38, 2193–2201.
- Huang, H.-P.; Sun, D.; Liu, Y.; Chu, W.-S.; Xiao, T.; Yuan, J.; Adam, H.; and Yang, M.-H. 2022. Adaptive transformers for robust few-shot cross-domain face anti-spoofing. In *European conference on computer vision*, 37–54. Springer.
- Huang, W.; Zheng, X.; Ma, X.; Qin, H.; Lv, C.; Chen, H.; Luo, J.; Qi, X.; Liu, X.; and Magno, M. 2024. An Empirical Study of LLaMA3 Quantization: From LLMs to MLLMs. arXiv:2404.14047.
- Jia, Y.; Zhang, J.; Shan, S.; and Chen, X. 2020. Single-side domain generalization for face anti-spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8484–8493.
- Jiang, D.; Liu, Y.; Liu, S.; Zhang, X.; Li, J.; Xiong, H.; and Tian, Q. 2023. From clip to dino: Visual encoders shout in multi-modal large language models.
- Jiang, Y.; Yan, X.; Ji, G.-P.; Fu, K.; Sun, M.; Xiong, H.; Fan, D.-P.; and Khan, F. S. 2024. Effectiveness assessment of recent large vision-language models. *Visual Intelligence*, 2(1): 17.
- Le, B. M.; and Woo, S. S. 2024. Gradient alignment for cross-domain face anti-spoofing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 188–199.
- Li, H.; Li, W.; Cao, H.; Wang, S.; Huang, F.; and Kot, A. C. 2018. Unsupervised domain adaptation for face anti-spoofing. *IEEE Transactions on Information Forensics and Security*, 13(7): 1794–1809.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Liao, C.-H.; Chen, W.-C.; Liu, H.-T.; Yeh, Y.-R.; Hu, M.-C.; and Chen, C.-S. 2023. Domain invariant vision transformer learning for face anti-spoofing. In *WACV*, 6098–6107.
- Liu, A. 2024. CA-MoEiT: Generalizable Face Anti-spoofing via Dual Cross-Attention and Semi-fixed Mixture-of-Expert. *International Journal of Computer Vision*, 1–14.
- Liu, A.; Ma, H.; Zheng, J.; Yuan, H.; Yu, X.; Liang, Y.; Escalera, S.; Wan, J.; and Lei, Z. 2024a. FM-CLIP: Flexible Modal CLIP for Face Anti-Spoofing. In *ACM MM*, 8228–8237.
- Liu, A.; Xue, S.; Gan, J.; Wan, J.; Liang, Y.; Deng, J.; Escalera, S.; and Lei, Z. 2024b. CFPL-FAS: Class Free Prompt Learning for Generalizable Face Anti-spoofing. In *CVPR*, 222–232.
- Liu, A.; Zhao, C.; Yu, Z.; Wan, J.; Su, A.; Liu, X.; Tan, Z.; Escalera, S.; Xing, J.; Liang, Y.; et al. 2022a. Contrastive context-aware learning for 3d high-fidelity mask face presentation attack detection. *TIFS*, 17: 2497–2507.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2024c. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Liu, S.; Lu, S.; Xu, H.; Yang, J.; Ding, S.; and Ma, L. 2022b. Feature generation and hypothesis verification for reliable face anti-spoofing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 1782–1791.
- Liu, S.-Q.; Lan, X.; and Yuen, P. C. 2018. Remote photoplethysmography correspondence feature for 3D mask face presentation attack detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 558–573.
- Liu, Y.; Chen, Y.; Dai, W.; Gou, M.; Huang, C.-T.; and Xiong, H. 2022c. Source-free domain adaptation with contrastive domain alignment and self-supervised exploration for face anti-spoofing. In *European Conference on Computer Vision*, 511–528. Springer.
- Liu, Y.; Chen, Y.; Dai, W.; Gou, M.; Huang, C.-T.; and Xiong, H. 2024d. Source-Free Domain Adaptation With Domain Generalized Pretraining for Face Anti-Spoofing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

- Liu, Y.; Chen, Y.; Gou, M.; Huang, C.-T.; Wang, Y.; Dai, W.; and Xiong, H. 2023. Towards unsupervised domain generalization for face anti-spoofing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 20654–20664.
- Liu, Y.; Jourabloo, A.; and Liu, X. 2018. Learning deep models for face anti-spoofing: Binary or auxiliary supervision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 389–398.
- 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.
- Shao, R.; Lan, X.; Li, J.; and Yuen, P. C. 2019. Multi-adversarial discriminative deep domain generalization for face presentation attack detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10023–10031.
- Shi, Y.; Gao, Y.; Lai, Y.; Wang, H.; Feng, J.; He, L.; Wan, J.; Chen, C.; Yu, Z.; and Cao, X. 2024. Shield: An evaluation benchmark for face spoofing and forgery detection with multimodal large language models. *arXiv preprint arXiv:2402.04178*.
- Srivatsan, K.; Naseer, M.; and Nandakumar, K. 2023. Flip: Cross-domain face anti-spoofing with language guidance. In *CVPR*, 19685–19696.
- Sun, Y.; Liu, Y.; Liu, X.; Li, Y.; and Chu, W.-S. 2023. Rethinking domain generalization for face anti-spoofing: Separability and alignment. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 24563–24574.
- Wang, C.-Y.; Lu, Y.-D.; Yang, S.-T.; and Lai, S.-H. 2022a. Patchnet: A simple face anti-spoofing framework via fine-grained patch recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20281–20290.
- Wang, J.; Zhang, J.; Bian, Y.; Cai, Y.; Wang, C.; and Pu, S. 2021. Self-domain adaptation for face anti-spoofing. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 2746–2754.
- Wang, K.; Zhang, G.; Yue, H.; Liang, Y.; Huang, M.; Zhang, G.; Han, J.; Ding, E.; and Wang, J. 2024a. CSDG-FAS: Closed-Space Domain Generalization for Face Anti-spoofing. *International Journal of Computer Vision*, 1–14.
- Wang, K.; Zhang, G.; Yue, H.; Liu, A.; Zhang, G.; Feng, H.; Han, J.; Ding, E.; and Wang, J. 2024b. Multi-domain incremental learning for face presentation attack detection. In *AAAI*, volume 38, 5499–5507.
- Wang, Z.; Wang, Z.; Yu, Z.; Deng, W.; Li, J.; Gao, T.; and Wang, Z. 2022b. 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.
- 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.
- 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.
- Yuan, Y.; Li, W.; Liu, J.; Tang, D.; Luo, X.; Qin, C.; Zhang, L.; and Zhu, J. 2024. Osprey: Pixel understanding with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28202–28211.
- Yue, H.; Wang, K.; Zhang, G.; Feng, H.; Han, J.; Ding, E.; and Wang, J. 2023. Cyclically disentangled feature translation for face anti-spoofing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 3358–3366.
- Zanella, L.; Menapace, W.; Mancini, M.; Wang, Y.; and Ricci, E. 2024. Harnessing Large Language Models for Training-free Video Anomaly Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18527–18536.
- Zhang, K.-Y.; Yao, T.; Zhang, J.; Liu, S.; Yin, B.; Ding, S.; and Li, J. 2021. Structure destruction and content combination for face anti-spoofing. In *2021 IEEE International Joint Conference on Biometrics (IJCB)*, 1–6. IEEE.
- Zhang, K.-Y.; Yao, T.; Zhang, J.; Tai, Y.; Ding, S.; Li, J.; Huang, F.; Song, H.; and Ma, L. 2020a. Face anti-spoofing via disentangled representation learning. In *ECCV*, 641–657. Springer.
- Zhang, S.; Roller, S.; Goyal, N.; Artetxe, M.; Chen, M.; Chen, S.; Dewan, C.; Diab, M.; Li, X.; Lin, X. V.; et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Zhang, Y.; Colman, B.; Guo, X.; Shahriyari, A.; and Bharaj, G. 2025. Common Sense Reasoning for Deepfake Detection. In *European Conference on Computer Vision*.
- Zhang, Y.; Yin, Z.; Li, Y.; Yin, G.; Yan, J.; Shao, J.; and Liu, Z. 2020b. Celeba-spoof: Large-scale face anti-spoofing dataset with rich annotations. In *ECCV*, 70–85.
- Zhang, Z.; Yan, J.; Liu, S.; Lei, Z.; Yi, D.; and Li, S. Z. 2012. A face antispoofing database with diverse attacks. In *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*, 175–187.
- 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 *CVPR*, 20453–20463.
- Zhou, Q.; Zhang, K.-Y.; Yao, T.; Yi, R.; Sheng, K.; Ding, S.; and Ma, L. 2022. Generative domain adaptation for face anti-spoofing. In *European Conference on Computer Vision*, 335–356. Springer.