

# Supervised Contrastive Learning for Snapshot Spectral Imaging Face Anti-Spoofing

Chuanbiao Song<sup>1</sup>, Yan Hong<sup>1</sup>, Jun Lan<sup>1</sup>,  
Huijia Zhu<sup>1</sup>, Weiqiang Wang<sup>1</sup>, Jianfu Zhang<sup>2\*</sup>

<sup>1</sup>Ant Group, <sup>2</sup>Qing Yuan Research Institute, Shanghai Jiao Tong University

<sup>1</sup>songchuanbiao.scb@antgroup.com, <sup>1</sup>yanhong.sjtu@gmail.com,

<sup>1</sup>{yelan.lj, huijia.zhj, weiqiang.wwq}@antgroup.com, <sup>2</sup>c.sis@sjtu.edu.cn

## Abstract

*This study reveals a cutting-edge re-balanced contrastive learning strategy aimed at strengthening face anti-spoofing capabilities within facial recognition systems, with a focus on countering the challenges posed by printed photos, and highly realistic silicone or latex masks. Leveraging the HySpeFAS dataset, which benefits from Snapshot Spectral Imaging technology to provide hyperspectral images, our approach harmonizes class-level contrastive learning with data resampling and an innovative real-face oriented reweighting technique. This method effectively mitigates dataset imbalances and reduces identity-related biases. Notably, our strategy achieved an unprecedented 0.0000% Average Classification Error Rate (ACER) on the HySpeFAS dataset, ranking first at the Chalearn Snapshot Spectral Imaging Face Anti-spoofing Challenge on CVPR 2024.*

## 1. Introduction

Face recognition technologies [1, 2, 9, 14, 15, 19, 22], with the extensive applications in various aspects of our life like mobile payments and access control systems, has significantly enhanced convenience. Nonetheless, its susceptibility to diverse forms of attacks limits its reliable application. Numerous malicious attacks, including the use of printed photos, video replays, and faces with flexible masks, can readily mislead these systems into making wrong judgments. To ensure the dependable operation of face recognition systems, face anti-spoofing (FAS) [10, 23, 25, 25] methods are crucial for identifying and mitigating various attacks.

Confronted with the challenge of highly convincing silicone or latex masks, the deployment of innovative spectroscopy sensors [17, 20] can notably boost the discriminative power of face recognition systems against these attacks.

Snapshot Spectral Imaging (SSI) [3, 7, 8] technologies possess the ability to capture compressed sensing spectral images, positioning it as an effective tool for the integration of spectroscopic information into current face recognition systems. Recently, utilizing a snapshot spectral camera, the Chalearn Snapshot Spectral Imaging Face Anti-spoofing Challenge at CVPR 2024 successfully acquires SSI images of both real and fake faces, and creates the first snapshot spectral face anti-spoofing dataset, named HySpeFAS. This dataset encompasses 6760 hyperspectral images, each reconstructed from SSI images using the TwIST [18] algorithm and featuring 30 spectral channels. These data present invaluable opportunities for FAS to advance the sophistication and reliability of algorithms.

In this paper, we present our approach tailored for the FAS task. We design a re-balanced contrastive learning approach, aimed at capturing the detailed and intrinsic patterns from the imbalanced dataset. We embed class-level contrastive learning into FAS task by employing data resampling to mitigate class-level imbalances in the dataset. Furthermore, we introduce an innovative real face-oriented reweighting methodology to effectively eliminate potential bias to the identity of the face. The proposed method achieves 0.0000% ACER on the HySpeFAS dataset and ranks first place on the Snapshot Spectral Imaging Face Anti-spoofing Challenge at CVPR 2024.

Data	Total	Fake Class	Real Class
Train	3900	3380	520
Val	936	728	208
Test	1924	-	-

Table 1. The split of train/validation/test real images and fake images on the HySpeFAS dataset.

\*Corresponding author.




Figure 1. Analysis of examples from HySpeFAS dataset. From top to bottom: real images and fake images. From left to right: examples in the first box show the same ID face across real and fake images; the second one indicates the various alternations of real and fake images; the third one visualizes different appearances of images from the same ID; the last one shows different orientations of real images and fake images.

## 2. HySpeFAS dataset

The HySpeFAS dataset utilizes a snapshot spectral camera to obtain SSI images of both real and fake faces, and those images are reconstructed from SSI images by TwIST algorithm and characterized with 30 spectral channels. In total, the dataset provides 6760 hyperspectral images, as detailed in Table 1. For the Snapshot Spectral Imaging Face Anti-spoofing Challenge at CVPR 2024, the organizers have provided all images along with their spectral matrices. The dataset is divided into a training set with 3900 images and a validation set containing 936 images. Visual examples from the dataset are shown in Figure 1, facilitating an in-depth analysis of the HySpeFAS dataset. This analysis serves as the foundation for data resampling and reweighting strategies discussed in Section 3.

From Table 1, it is evident that although the dataset is valuable due to the challenging acquisition process, the quantity of data remains limited. Additionally, the number of counterfeit face images significantly outnumbers that of genuine faces, creating a pronounced imbalance between the two primary classes of images. In Figure 1, several key characteristics of the HySpeFAS dataset are identified: (1). identical identifiers (ID) are present in both fake and real face images. (2). face images exhibit a variety of alterations, including masks and transparent masks. (3). faces from the same ID show considerable variation in physical appearance. (4). variations in facial orientation and other conditions are also observed. These characteristics present considerable challenges, and the methodology proposed in this paper is designed to address these specific aspects.

## 3. Methodology

In this section, we first present the data preprocessing for the HySpeFAS dataset on the basis of data analysis in Section 2. Then we introduce our framework comprised of multiple modules. Next, we describe the used loss function compatible with different modules. Finally, the intra-class

*mixup* [26], the real face-oriented reweighting, and the cross batch memory [21] are integrated into the training strategy to promote the supervised contrastive learning.

### 3.1. Data Preprocessing

Given the volume and the imbalance between the real and fake sample quantities within the dataset, we initiate our approach with preprocessing enhancements to address this issue.

**Class Balancing** As shown in Table 1, we analyze the class numbers of the training data and validation data, and can observe that the HySpeFAS dataset is an unbalanced dataset, where the number of fake data is much larger than that of real data. To eliminate the effects of the imbalanced data, we adopt the oversampling strategy to rebalance the data distribution by amplifying the volume of real instances.

**Data Augmentations** During the training, we use extensive data augmentations, such as *random crop*, *random horizontal flip*, *cutout* [5] and *random mask*. For *random mask*, we randomly mask the bottom half of training samples to eliminate the effect of the worn mask, or randomly mask the left or right half of training samples. Figure 3 presents some examples of augmented faces based on *random mask*.

### 3.2. Framework

Our framework, shown in Figure 2, leverage the multi-attention network MAT [27] as the backbone  $\mathbf{E}_{mat}$ , and combine the spectral weight learning module [6]  $\mathbf{E}_{swl}$ , the central differential convolution [24]  $\mathbf{E}_{cdc}$ , the classifier  $\mathbf{E}_c$ , and the contrastive learning module  $\mathbf{E}_{scl}$ . image-label pairs from HySpeFAS dataset can be represented as  $\{\mathbf{x}_I, \mathbf{x}_m, \mathbf{y}\}$ , where  $\mathbf{x}_I \in \mathbb{R}^{w \times h \times 3}$  (*resp.*,  $\mathbf{x}_m \in \mathbb{R}^{w \times h \times 30}$ , and  $\mathbf{y} \in \mathbb{R}^2$ ) denotes RGB image (*resp.*, spectral matrix, and one-hot label). We concatenate  $\mathbf{x}_I$  and  $\mathbf{x}_m$  along the third channel to form input sample  $\mathbf{x} \in \mathbb{R}^{w \times h \times 33}$ .

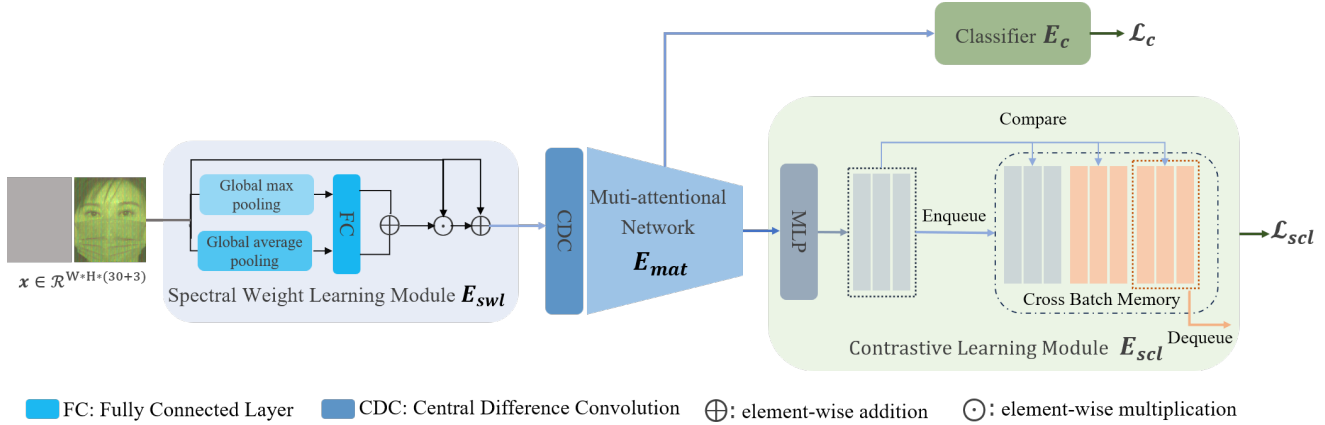


Figure 2. The framework of our proposed method.




Figure 3. Visualization of augmented examples with three different type of *random mask*.

Given the multimodal nature of the input, input sample  $x$  first is processed with the spectral weight learning module  $E_{swl}$  to assign different weights to input channels. Then we adopt the  $E_{cdc}$  as the first convolution layer to obtain gradient-level features, and use a multi-attentional network  $E_{mat}$  to learn discriminative features. Feded with the learned discriminative features, the classifier  $E_c$  is responsible for distinguishing fake faces from real faces, while the contrastive learning module  $E_{scl}$  is dedicated to promoting the discriminability of the learned features.

### 3.3. Training Strategies

#### 3.3.1 Intra-class Mixup

Given the limited size of the training dataset, the risk of overfitting is heightened when employing large neural networks. To mitigate this and enrich the diversity of training samples, we utilize a variant of *mixup*, inter-class *mixup*. As formulated by Eqn. 1, inter-class *mixup* generates training samples  $(\hat{x}, \hat{y})$  by interpolating between two training samples from the same class.

$$\begin{aligned} \hat{x} &= \lambda \cdot x_i + (1 - \lambda) \cdot x_j, \\ \hat{y} &= \lambda \cdot y_i + (1 - \lambda) \cdot y_j, \end{aligned} \quad (1)$$

where  $y_i$  and  $y_j$  are the one-hot labels of the same class, and  $\lambda \in [0, 1]$  represents the mixing parameter. Following the setting of the original *mixup*, we set  $\lambda \sim \text{Beta}(1.0, 1.0)$ .

#### 3.3.2 Real-face Oriented Reweighting

To diminish the model’s reliance on content irrelevant to spoofing detection, such as identity and facial features, we introduce Real-face Oriented Reweighting (ROR) strategy during training. ROR assigns weights to fake training samples based on their face cosine similarity with real training samples. The face cosine similarities are calculated based on the typical face model ArcFace [4] as follows,

$$w_{x_i^F} = \max \frac{1 + \cos(f_{\text{Arcface}}(x_i^F), f_{\text{Arcface}}(x_j^R))}{2}, \quad (2)$$

where  $x_i^R$  (*resp.*,  $x_j^F$ ) denotes the real face image (*resp.*, fake face images). Fake samples exhibiting higher similarity to real faces are assigned greater weights. This reweighting approach encourages the model to deprioritize learning from features strongly tied to identity and facial appearance, thereby focusing more on distinguishing genuine from faces.

#### 3.3.3 Objective Functions

We design two distinct loss functions: one loss for classification, and a contrastive loss for regularizing real and fake data features. These losses are combined in a weighted sum to create the overall loss function for training the framework.

**Focal Loss.** Instead of using the typical cross entropy for classification, we adopt the focal loss [12], which is based

on a variant of cross entropy for binary classification:

$$\mathcal{L}_c = \text{FL}(p_t) = -(1 - p_t)^\gamma \log(p_t), \quad (3)$$

where  $p_t$  is the probability that the model predicts for the ground truth object and we set  $\gamma$  as 2. According to Eqn. 3, the focal loss gives less weight to easy examples and gives more weight to hard misclassified examples.

**Supervised Contrastive Loss.** The objective of contrastive regularization loss is to optimize the similarity and dissimilarity of real and fake data embeddings. The contrastive regulation loss is formulated as:

$$\mathcal{L}_i^{\text{sup}} = \sum_{j=1}^N 1_{i \neq j} \cdot 1_{\tilde{y}_i = \tilde{y}_j} \cdot \log \frac{\exp(z_i \cdot z_j / \tau)}{\sum_{k=1}^N 1_{i \neq k} \cdot \exp(z_i \cdot z_k / \tau)},$$

$$\mathcal{L}_{\text{scl}} = - \sum_{i=1}^N \mathcal{L}_i^{\text{sup}}, \quad (4)$$

where  $z_i$  is denoted as the normalized embedding of the training sample  $x_i$  from  $\mathbf{E}_{\text{scl}}$  module, and  $\tau$  serves as a temperature hyper-parameter.

According to Eqn. 4, the supervised contrastive loss maximizes the cosine similarity between the training samples with the same category, while simultaneously minimizing the cosine similarity between the training samples with different categories. We compute the loss between real and fake samples to encourage the model to learn a generalizable representation that across different images.

To boost the performance of the supervised contrastive loss, we utilize the cross batch memory (XBM) [21] to collect sufficient hard negative pairs for contrastive learning. Specifically, XBM memorizes the embeddings of recent mini-batches and can provide sufficient embeddings for calculating the supervised contrastive loss. It operates on a queue principle, enrolling the latest batch of embeddings while simultaneously removing the oldest, maintaining a dynamic and up-to-date memory bank for optimization.

**Overall Loss.** The final loss function of the training process is the weighted sum of the above loss functions:

$$\mathcal{L} = \mathcal{L}_c + \lambda_{\text{scl}} \cdot \mathcal{L}_{\text{scl}}, \quad (5)$$

where  $\lambda_{\text{scl}}$  is a hyper-parameter for balancing the overall loss, and we set  $\lambda_{\text{scl}}$  as 10 for a strong regularization.

## 4. Experiments

In this section, we will describe the evaluation metrics, training details, as well as the performance of our proposed method on the HySpeFAS dataset.

### 4.1. Evaluation Metrics

Following the HySpeFAS dataset, we select the Attack Presentation Classification Error Rate (APCER), Bona Fide Presentation Classification Error Rate (BPCER), and Average Classification Error Rate (ACER) as the evaluation metric. Specifically, APCER and BPCER are formulated as below:

$$APCER = \frac{FN}{TP + FN}, \quad BPCER = \frac{FP}{FP + TN}, \quad (6)$$

where FN(False Negative) and FP(False Positive) refer to the number of incorrectly classified fake or real samples respectively, and TP(True Positive) and TN(True Negative) refer to the number of correctly classified real or fake samples respectively. The ACER is calculated as below:

$$ACER = \frac{APCER + BPCER}{2}, \quad (7)$$

which is used as the main evaluation metric on the test set and determine the final ranking of the competition. The lower the ACER value, the better the performance.

### 4.2. Training Details

We implement our method on 1 NVIDIA Tesla A100 80G GPU based on open-source framework PyTorch [16]. We train the network using the ASAM optimizer [11] with 30 epochs. The learning rate is 0.01 initially and adjusted by the cosine annealing schedule [13]. The batch size is 240 and the weight decay is  $5 \times 10^{-3}$ , and the memory size of the XBM [21] is 1200. The temperature parameter  $\tau$  of the supervised contrastive loss is set to 0.07.

### 4.3. Performance Results

We compare the performance of our method and the solutions of other teams on the test set in the Snapshot Spectral Imaging Face Anti-spoofing Challenge at CVPR 2024. The evaluation scores of ours and other teams are shown in Table 2. We can observe that all the top 10 teams achieve excellent performance results, where all the ACER results are less than 1%. Our method achieves ACER, APCER, and BPCER by 0%, 0% and 0%, respectively, ranking the first place in this competition.

## 5. Conclusion

In this paper, we introduce supervised contrastive learning for snapshot spectral imaging face anti-spoofing based on the multi-attention neural network. Furthermore, to boost the supervised contrastive learning, we utilize the intra-class mixup to improve the diversity of training samples, the real-face oriented sample reweighting to avoid the effects of the identity feature, and the cross-batch memory to



Team	ACER(%)	APCER(%)	BPCER(%)
DXAI	0.7237	0.8065	0.6410
ZTT	0.6927	0.7444	0.6410
galileo	0.6359	0.3102	0.9615
kk_li	0.6359	0.3102	0.9615
ctyun-ai	0.4756	0.3102	0.6410
Ricardozzf	0.2223	0.1241	0.3205
hexianhua	0.1861	0.3722	0.0000
SeaRecluse	0.0310	0.0620	0.0000
CTEL_AI	0.0000	0.0000	0.0000
<b>Ours</b>	<b>0.0000</b>	<b>0.0000</b>	<b>0.0000</b>

Table 2. The top10 leaderboard of the Snapshot Spectral Imaging Face Anti-spoofing Challenge at CVPR 2024.

increase the number of the contrastive samples. Experimental results show that the proposed method achieves excellent performance and yields the first place among all teams on the recently conducted the Snapshot Spectral Imaging Face Anti-spoofing Challenge at CVPR 2024.

## References

- [1] Gwangbin Bae, Martin de La Gorce, Tadas Baltrušaitis, Charlie Hewitt, Dong Chen, Julien Valentin, Roberto Cipolla, and Jingjing Shen. Digiface-1m: 1 million digital face images for face recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3526–3535, 2023. 1
- [2] Fadi Boutros, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Elasticface: Elastic margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1578–1587, 2022. 1
- [3] Xun Cao, Tao Yue, Xing Lin, Stephen Lin, Xin Yuan, Qionghai Dai, Lawrence Carin, and David J Brady. Computational snapshot multispectral cameras: Toward dynamic capture of the spectral world. *IEEE Signal Processing Magazine*, 33(5):95–108, 2016. 1
- [4] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4690–4699. Computer Vision Foundation / IEEE, 2019. 3
- [5] Terrance Devries and Graham W. Taylor. Improved regularization of convolutional neural networks with cutout. *CoRR*, abs/1708.04552, 2017. 2
- [6] Ying Fu, Hongrong Liu, Yunhao Zou, Shuai Wang, Zhongxiang Li, and Dezhi Zheng. Category-level band learning-based feature extraction for hyperspectral image classification. *IEEE Trans. Geosci. Remote. Sens.*, 62:1–16, 2024. 2
- [7] Nathan Hagen and Michael W Kudenov. Review of snapshot spectral imaging technologies. *Optical Engineering*, 52(9):090901–090901, 2013. 1
- [8] Robert T Kester, Noah Bedard, Liang Gao, and Tomasz S Tkaczyk. Real-time snapshot hyperspectral imaging endo-scope. *Journal of biomedical optics*, 16(5):056005–056005, 2011. 1
- [9] Minchul Kim, Anil K Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18750–18759, 2022. 1
- [10] Jukka Komulainen, Abdenour Hadid, and Matti Pietikäinen. Context based face anti-spoofing. In *2013 IEEE Sixth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–8. IEEE, 2013. 1
- [11] Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. ASAM: adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, pages 5905–5914. PMLR, 2021. 4
- [12] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2999–3007. IEEE Computer Society, 2017. 3
- [13] Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with warm restarts. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. 4
- [14] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. Magface: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021. 1
- [15] Omkar Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In *BMVC 2015-Proceedings of the British Machine Vision Conference 2015*. British Machine Vision Association, 2015. 1
- [16] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 4
- [17] Shijie Rao, Yidong Huang, Kaiyu Cui, and Yali Li. Anti-spoofing face recognition using a metasurface-based snapshot hyperspectral image sensor. *Optica*, 9(11):1253–1259, 2022. 1
- [18] Hoover Rueda-Chacon, Juan F Florez-Ospina, Daniel L Lau, and Gonzalo R Arce. Snapshot compressive tof+ spectral imaging via optimized color-coded apertures. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2346–2360, 2019. 1
- [19] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. 1
- [20] Holger Steiner, Andreas Kolb, and Norbert Jung. Reliable face anti-spoofing using multispectral swir imaging. In *2016 international conference on biometrics (ICB)*, pages 1–8. IEEE, 2016. 1
- [21] Xun Wang, Haozhi Zhang, Weilin Huang, and Matthew R. Scott. Cross-batch memory for embedding learning. In

- 2020 *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 6387–6396. Computer Vision Foundation / IEEE, 2020. [2](#), [4](#)
- [22] Haiyu Wu, Vítor Albiero, KS Krishnapriya, Michael C King, and Kevin W Bowyer. Face recognition accuracy across demographics: Shining a light into the problem. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1041–1050, 2023. [1](#)
- [23] Jianwei Yang, Zhen Lei, and Stan Z Li. Learn convolutional neural network for face anti-spoofing. *arXiv preprint arXiv:1408.5601*, 2014. [1](#)
- [24] Zitong Yu, Chenxu Zhao, Zezheng Wang, Yunxiao Qin, Zhuo Su, Xiaobai Li, Feng Zhou, and Guoying Zhao. Searching central difference convolutional networks for face anti-spoofing. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 5294–5304. Computer Vision Foundation / IEEE, 2020. [2](#)
- [25] Zitong Yu, Chenxu Zhao, Zezheng Wang, Yunxiao Qin, Zhuo Su, Xiaobai Li, Feng Zhou, and Guoying Zhao. Searching central difference convolutional networks for face anti-spoofing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5295–5305, 2020. [1](#)
- [26] Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. [2](#)
- [27] Hanqing Zhao, Wenbo Zhou, Dongdong Chen, Tianyi Wei, Weiming Zhang, and Nenghai Yu. Multi-attentional deepfake detection. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 2185–2194. Computer Vision Foundation / IEEE, 2021. [2](#)