

Surveillance Face Presentation Attack Detection Challenge

Hao Fang¹, Ajian Liu², Jun Wan^{1,2*}, Sergio Escalera³, Hugo Jair Escalante⁴, Zhen Lei^{1,2,5}

¹School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

²MAIS, Institute of Automation, Chinese Academy of Sciences, Beijing, China

³Universitat de Barcelona and Computer Vision Center, Barcelona, Catalonia, Spain

⁴Instituto Nacional de Astrofísica, Óptica y Electrónica, Puebla, Mexico

⁵CAIR, HKISI, CAS

fanghao21@mails.uicas.ac.cn, ajianliu92@gmail.com, jun.wan@ia.ac.cn

Abstract

Face Anti-spoofing (FAS) is essential to secure face recognition systems from various physical attacks. However, most of the studies lacked consideration of long-distance scenarios. Specifically, compared with FAS in traditional scenes such as phone unlocking, face payment, and self-service security inspection, FAS in long-distance such as station squares, parks, and self-service supermarkets are equally important, but it has not been sufficiently explored yet. In order to fill this gap in the FAS community, we collect a large-scale Surveillance High-Fidelity Mask (SuHi-FiMask). SuHiFiMask contains 10,195 videos from 101 subjects of different age groups, which are collected by 7 mainstream surveillance cameras. Based on this dataset and protocol-3 for evaluating the robustness of the algorithm under quality changes, we organized a face presentation attack detection challenge in surveillance scenarios. It attracted 180 teams for the development phase with a total of 37 teams qualifying for the final round. The organization team re-verified and re-ran the submitted code and used the results as the final ranking. In this paper, we present an overview of the challenge, including an introduction to the dataset used, the definition of the protocol, the evaluation metrics, and the announcement of the competition results. Finally, we present the top-ranked algorithms and the research ideas provided by the competition for attack detection in long-range surveillance scenarios.

1. Introduction

In recent years, as face recognition systems are widely used, the security of face recognition systems has been increasingly threatened. Therefore, face Presentation Attack Detection (PAD) techniques are crucial in defending

against malicious attacks and securing face recognition systems. With the release of some high-quality 2D attack datasets [2, 6, 27, 36, 54–56] and 3D mask datasets [10, 14, 32, 33, 40, 41, 43], existing work [5, 13, 21, 26, 28, 30, 37, 42, 49, 51, 53] has achieved satisfactory performance in short-distance applications. However, these algorithms rely heavily on the quality of face images. As a result, their performance degrades significantly in long-distance scenarios, hindering the expansion of FAS work to surveillance contexts. With the popularity of remote cameras and the deployment of surveillance networks, face recognition technology [19, 20, 57] is gradually expanding to long-distance applications, such as real-time capture, identity verification in security channels, and payment in self-service supermarkets. Therefore, in order to safeguard society and citizens, we require vision technology that can effectively operate in long-distance surveillance scenarios.

However, the FAS community has conducted little research on deceptive face detection in long-distance surveillance environments. We identify two reasons hindering the development of PAD technology: (1) Lack of surveillance scenario-based datasets. Due to the difficulty of building surveillance scenes and expensive acquisition costs, the existing FAS datasets require subjects to face the acquisition device at a certain distance. However, diverse surveillance scenes, realistic weather and lighting, and natural human behavior are important data features of the remote surveillance PAD dataset. (2) Lack of a common benchmark based on remote surveillance dataset to compare the performance of different algorithms. Current research in the FAS community, both algorithms based on color texture feature learning [14, 17, 18, 50] and remote photoplethysmography (rPPG)-based detection [23, 34, 35] require high quality face images to extract fine-grained features. However, data based on long-distance surveillance scenarios bring new challenges to FAS algorithms such as low res-

*Contact person

olution, motion blur, and occlusion. Therefore, a common benchmark based on remote surveillance datasets to drive algorithm advancement is critical.

In order to facilitate the FAS community’s research on attack detection in long-distance surveillance scenarios, we targeted the following three aspects to address the current difficulties based on the above analysis: (1) We collected a large-scale PAD dataset based on surveillance scenarios, namely SuHiFiMask [11]. Compared to existing short-distance datasets, it has several advantages, such as diverse remote surveillance scenarios, realistic distribution of human faces, a rich pool of spoofing attacks, realistic illumination and diverse weather. (2) We define a more general test protocol for real-world FAS algorithm deployment and publish a public benchmark based on remote surveillance scenarios. We design protocols to evaluate the generality of the algorithms for ID changes, mask type changes, and picture quality changes. (3) Based on the dataset and protocol, we successfully held the *Surveillance Face Presentation Attack Detection Challenge at CVPR2023*¹, which attracted 180 teams from all over the world and significantly boosted the community’s research on FAS tasks in surveillance scenarios. A summary of the team names and affiliations that made it to the final stage is shown in Tab 1. We found that most of the teams that qualified for the final round of this competition used large networks as the backbone for their algorithms compared to the previous challenge [1, 25, 29, 31], suggesting that large models are gradually outperforming traditional vision models on downstream tasks.

Table 1. Team and affiliations are listed in the final ranking of this challenge.

Ranking	Team Name	Leader Name, Affiliation
1	MateoH	Keyao Wang, Baidu Inc
2	CTEL_AI	Yaowen Xu, China Telecom
3	horsego	Dingheng Zeng, MaShang-Consumer Finance Co., Ltd
4	hexianhua	Xianhua He, Meituan
5	OPDAI	Heng Cong, NetEase Inc
6	SeaRecluse	Minzhe huang, akuvox
7	XiangR	Xiang Rao, Xiamen University
8	Chenyifan	TengTeng Zhang, CMB
9	ioNetworks	Yingyu Chen, ioNetworks

2. Challenge Overview

In this section, we review the organized challenge, including a brief introduction to the surveillance-based dataset SuHiFiMask, an introduction to the protocol produced for this challenge, the process and timeline of the

¹<https://sites.google.com/view/face-anti-spoofing-challenge/welcome/challengecvpr2023?authuser=0>

challenge, and the evaluation metrics.

2.1. SuHiFiMask Dataset

To the best of our knowledge, SuHiFiMask is the first dataset to extend FAS to real surveillance scenes rather than mimicking low-resolution images and surveillance environments. It contains 10,195 videos, which were taken from 101 subjects of different age groups. For the attack pool, it contains 232 3D high-fidelity masks (plaster, resin, silicone, and head mold), 200 2D attacks (posters, portraits, and screens), and 2 adversarial attacks (adversarial hats, adversarial masks). During the collection process, it considers 40 surveillance scenes for video recording, including daily life scenes (e.g., cafes, cinemas, and theaters) and security check scenes (e.g., security check lanes and parking lots) for deploying face recognition systems. In addition, SuHiFiMask takes into account four different types of weather (Sunny, Windy, Cloudy, and Snowy days) and natural lighting (Day and Night light) to imitate the complex and changing real environment. Finally, seven mainstream surveillance cameras were used to record video simultaneously to ensure complete capture of the entire scene.

In order to facilitate the use of the dataset by the participating teams, we performed the following pre-processing: (1) We used RetinaFace [8] to detect the face in each frame of the original video. (2) We use face similarity to track the position of faces in consecutive frames. (3) We sample each video in 10-frame intervals and store the cropped face image in the corresponding face tracking box folder. (4) We named the folder of this video according to the following rule: *Group_Scene_Camera_Epoch_Time*.

2.2. Challenge Protocol and Data Statistics

In order to promote the extension of the FAS task to long-range scenarios and to meet the deployment requirements of realistic surveillance networks, we designed a protocol that can evaluate the generalization ability of the algorithm under image quality variations. The inference phase of the algorithm degrades in performance due to the input of low-quality images containing noise, which indicates the inadequate generalization ability of the algorithm. Therefore, to better evaluate the generalization ability of the algorithm under quality variations, we constitute the training set, development set and testing set with images of different quality distributions. In detail, we use SER-FIQ [45] algorithm to calculate the image quality score which ranges from 0 to 1. As shown in the fifth column of Tab. 2, we assigned images with scores [0.4, 1] as the training set, scores [0.3, 0.4] as the development set, and scores [0, 0.3] as the testing set. In addition, to accurately evaluate the robustness of the algorithm to quality changes without being affected by other factors, all mask types (referred to as 1,2,3,4), subjects, scenes, weather (Sunny, Windy, Cloudy and Snowy

days), and lighting (Day and Night light) were included in the training set, the development set, and the testing set.

Table 2. Statistical information for each protocol of the proposed SuHiFiMask dataset. Note that 1, 2, and 3 in the fourth column mean resin, silicone, and plaster. 4 represents headgear and head mold.

Prot.	SubSet	#Subj.	Mask	Score	#Live	#Mask	#Others	#All
3	Train	101	1~4	[0.4,1]	64,276	35,898	58,889	159,063
	Dev	101	1~4	[0.3,0.4]	37,990	24,031	27,255	89,276
	Test	101	1~4	[0,0.3]	83,181	43,087	35,614	161882

2.3. Challenge Process and Timeline

The challenge was run in the CodaLab² platform, and this section will introduce the two stages of this challenge.

2.3.1 Development Phase

(Started: February. 1, 2023 - Ended: February 24, 2023) During this phase, teams have access to labeled training data and unlabeled development data. Participants could use the training data to train their models and submit predictions on the development data. The training data and development data contain all subjects, attack types, surveillance scenarios, weather, lights, and sensors, the only difference is the quality of the images in the training and development sets. In addition, labels for the development set are not provided to participants, who could submit predictions for development data and receive real-time ranking feedback via the leaderboard.

2.3.2 Final Phase

(Started: February. 24, 2023 - Ended: March 5, 2023) In the final phase, the labels of the development set are provided to the teams so that the teams incorporate the development data into the training phase of the model. At the same time, testing sets without labels were published and participants had to make predictions on the test sets and upload the results of their algorithms to the challenge platform. During the final phase, teams will have two opportunities per day to submit their solutions to the challenge platform for the purpose of assessing the stability of their algorithms. Note that the CodaLab platform will display the results of the last submission by default on the challenge website. In addition, to ensure that the final ranking is authentic, teams must publicly release their code under the license of their choice and provide a fact sheet describing their solution.

²<https://codalab.lisn.upsaclay.fr/competitions/10080>

2.4. Evaluation Metrics

For the performance evaluation, we selected the recently standardized ISO/ IEC 30107-3 metrics: Attack Presentation Classification Error Rate (APCER), Normal/ Bona Fide Presentation Classification Error Rate (NPCER/ BPCER) and Average Classification Error Rate (ACER) as the evaluation metric, in which APCER and BPCER/ NPCER are used to measure the error rate of fake or live samples, respectively. The ACER on the testing set is determined by the Equal Error Rate (EER) thresholds on the development set. In addition, the Area Under Curve (AUC) is defined as the area under the ROC curve and its value generally ranges from 0.5 to 1. In this competition, AUC was used as an additional evaluation criterion. Finally, the final ranking of this challenge was determined by the value ACER.

3. Description of solutions

3.1. MateoH

Considering that there are many indistinguishable samples in the protocol, the team of MateoH proposes a progressive training strategy (PTS) to solve the catastrophic forgetting problem brought by non-stationary data to the deep network. As shown in Fig. 1, the process of PTS is as follows: (1) Sample the initial set of samples. PTS will randomly select a balanced number of positive and negative samples from the training set as the initial training sample set based on the initial sampling rate. The remaining samples are put into the pending set for validation. (2) Hard sample mining. After each training round, the samples in the pending set are predicted with the model and ranked according to their prediction scores. Positive samples with a low score and negative samples with a high score will be added to the training set according to the sampling rate. (3) Reduced sampling rate. A decay factor is used to reduce the sampling rate after each round as a way to gradually increase the focus on hard samples. Overall, the PTS training strategy provides a framework paradigm for progressive hard-sample mining. The strategy allows the model to progressively strengthen its focus on hard samples while playing back previous simple samples to avoid catastrophic forgetting.

To address the issue of the limited defense capability of FAS systems against unknown attack types in surveillance scenarios, the team of MateoH proposes a Dynamic Feature Queue (DFQ) module. This module views the negative sample set as a cluster of multiple unknown categories and the positive sample set as a closed sample set. Thus, the binary classification task in FAS is modeled as a multi-classification task thus improving the model's detection of multiple attack types. The training flow of this module is shown in Fig. 1: (1) ViT-Large is used as a backbone network to extract image features. (2) Calculate the cosine similarity \log_0 of the feature to the negative sample center and

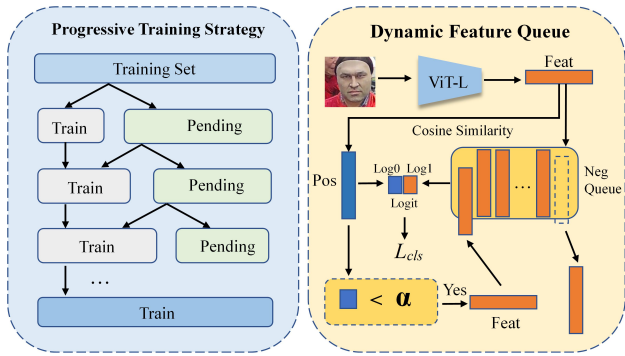


Figure 1. Progressive Training Strategy (PTS) and Dynamic Feature Queue (DFQ) module proposed by team of MateoH. The left side of the figure shows that PTS divides the original training set into a training subset and a pending set, and gradually adds hard samples from the pending set to the training subset. The right side of the figure shows the training process of the DFQ module.

the highest cosine similarity \log_1 of the feature to the negative sample center queue. (3) Concatenate \log_0 and \log_1 as logit and calculate the cross entropy loss of logit and label. (4) Judge the size of \log_0 and threshold α , and add the negative samples with \log_0 less than the threshold to the negative sample queue. (5) Remove the queue head element from the queue and update the network parameters.

3.2. CTEL_AI

Since the training, development, and testing sets for this challenge have different image quality distributions, models with insufficient generalization capabilities will experience degraded performance under the influence of quality differences. To address this problem, the CTEL_AI team proposes an adversarial domain generalization method that learns a consistent distribution of deceptive features to avoid models that rely too much on detailed features of high-quality images. Specifically, they delineate the domain using the quality score as the boundary, and consider the training set with the picture score interval of $[0.4, 1]$ as domain one and the development set with the picture quality score interval of $[0.3, 0.4)$ as domain two. By optimizing the adversarial domain loss, the encoder learns to extract consistent quality-independent deceptive features. As shown in Fig. 2, the CTEL_AI team uses the ViT-Large model [9] as the encoder of the solution and adds the gradient reversal layer GRL [12] after the encoder to invert the gradient, thus allowing the parameters of the domain classifier to be optimized in the opposite direction.

In terms of training strategies, SGD was used for gradient descent with a momentum of 0.9 and an initial learning rate (lr) of 0.01. Cosine learning rate decay was employed and the lowest lr is the 1% of initial lr. An epoch was warm-up with a learning rate, and a total of 100 epochs




Figure 2. The framework of team CTEL_AI. The training set and the development set are divided into two different domains. After that, ViT-Large model is used to extract the features of the input images. Cross-entropy loss and adversarial loss are used to optimize the encoder to guide the encoder to extract quality consistent features.

were trained. It is worth mentioning that based on the above training strategy, they fine-tune the model for a total of 30 epochs after completing 100 epochs of training. In the fine-tuning phase, they used different data augmentation methods to expand the amount of data and gradually added it to the training set. The data augmentation methods include Coarse dropout on the image, Random flip, Random rotate, Random crop, Split patch and shuffle, Gray-scale spoof image, Motion blur, Gaussian blur, Sharpen, Down-sampling, Randomly erase and repeat patches, Random fog, Posterize. The data enhanced by different methods are gradually added to the training set to avoid conflicts caused by using too many enhancement methods at the same time, and also to verify the effectiveness of each method.

Finally, the loss of the whole scheme consists of the cross-entropy loss of the main classifier and the adversarial loss of the domain classifier, as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{cls} + \mathcal{L}_{adv} \quad (1)$$

3.3. horsego

Compared to high-quality images, the training set images of this competition have low resolution, and thus there is little discriminative information in the images for the network to learn. To solve this problem, the team of horsego filtered the original images through a band-pass filter to obtain additional frequency domain information. In addition, they concatenate the spatial domain feature and frequency domain feature of the images before the fully connected layer to fuse them across channels, which makes the binary classification results more accurate. As shown in Fig. 3, they use a dual-stream structure, where one branch performs data enhancement and transformation on the original image to obtain a three-channel image of size 224×224 . The other branch uses a band-pass filter to mask some high-frequency information and low-frequency information of the original image, and then converts the image into a filtered map of size 224×224 by inverse fast Fourier transform. Then, two

EfficientFormerV2 [22] are used as encoders to extract frequency domain features and spatial domain features of the processed images, respectively. Finally, these two features are fused across channels and fed to the main classifier for classification, and the cross-entropy loss is used to supervise the main classifier.




Figure 3. The framework of the team horsego. The AI and BI are obtained by data enhancement and band-pass filtering of the original image I. The two encoders then extracted the features of AI and BI separately and concatenated the features. Finally, the fused features are sent to the main classifier and the cross-entropy loss is used for the optimization of the network.

In detail, the steps for data pre-processing through band-pass filters are as follows: (a) The original image undergoes Fast Fourier Transform (FFT) transformation and translation to obtain a frequency spectrum F with energy centralized. (b) Two Gaussian kernels with different variances are subtracted to obtain the band-pass kernel K. (c) The point product of F and K is calculated to obtain the frequency domain band-pass filtered frequency spectrum BF. (d) BF undergoes translation, IFFT (Inverse Fast Fourier Transform) transformation, and is taken as a real value to obtain the band-pass filtering image (BI). In addition, Google’s newly proposed LION optimizer [4] is used for network training, which has faster convergence and higher accuracy compared to other optimizers.

3.4. hexianhua

In the acquisition process of the challenge dataset, multiple surveillance cameras capture the same face simultaneously from different distances and angles. The team of hexianhua refers to such images with different perspectives of the same identity as homologous images. They propose a new data pre-processing strategy to effectively simulate homologous images of different quality to expand the training set, so that the model learns features of different quality distributions. As shown in Fig. 4, in the RandomZoomInOut module, the original images are randomly reduced to between 0.20x and 0.75x, and the interpolation algorithms used for the reduction (Nearest neighbor interpolation, Linear interpolation, Area interpolation, Lanczos interpolation) are also chosen randomly. Then, the interpolation algorithms are again randomly selected to resize these low-

resolution images to 224×224 to fit the input size of the encoder.

Moreover, considering that the number of live samples and the number of attack samples in the training samples are not balanced, the team of hexianhua uses the data balancing strategy of Max Upsampling to mitigate the performance degradation caused by the long tail problem. As shown in the right side of Fig. 4, they use the total number of categories with more samples as the number of Up-sampling. Then they repeatedly selected images from the set of images in the fewer categories to add to the training set until the number of images in the two categories in the training set is equal. In addition, they use the training method of EMA (Exponential Moving Average) and the data post-processing method of TTA (Test Time Augmentation) to further improve the performance of the model.

Finally, SwinV2-Huge [38] is used as the backbone network of the method, and both cross-entropy loss and Focal loss [24] are used to constrain the encoder, making it give more attention to some hard samples while learning discriminative features. The total loss function formulates as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{cls} + 0.5 * \mathcal{L}_{focal} \quad (2)$$

3.5. OPDAI

The team of OPDAI proposed a dual-stream network, with ConvNextv2_base and ConvNeXtv2_large [48] as the backbone of the two branches. They also added an MLP for each branch to improve the representativeness of the features extracted from different branches by optimizing the respective classifier parameters. Finally, they concatenate the features extracted by the two extractors and classify the concatenated features with an independent MLP, which makes the features learned by the two encoders more effective after fusion. As shown in Fig. 5, three MLP layers are supervised with focal loss to enhance the model’s focus on difficult samples in surveillance scenarios. Where loss1 is the focal loss of MLP1 for the classification of fused features, and loss2 and loss3 are the focal losses of MLP2 and MLP3 for the classification of branching features, respectively. The total loss function is expressed as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{focal1} + 0.5 * \mathcal{L}_{focal2} + 0.5 * \mathcal{L}_{focal3} \quad (3)$$

In terms of training strategy, they contribute a scheme to optimize the overall network framework. First, they fixed the parameters of the two feature extractors and pre-trained the three MLPs with a batch size of 512. Then they fine-tuned the whole dual-stream network with a batch size of 200, and the training process continued until the sum of the losses dropped to 0.0001. This has the following two advantages: (1) Allowing the fully connected layer in the MLP

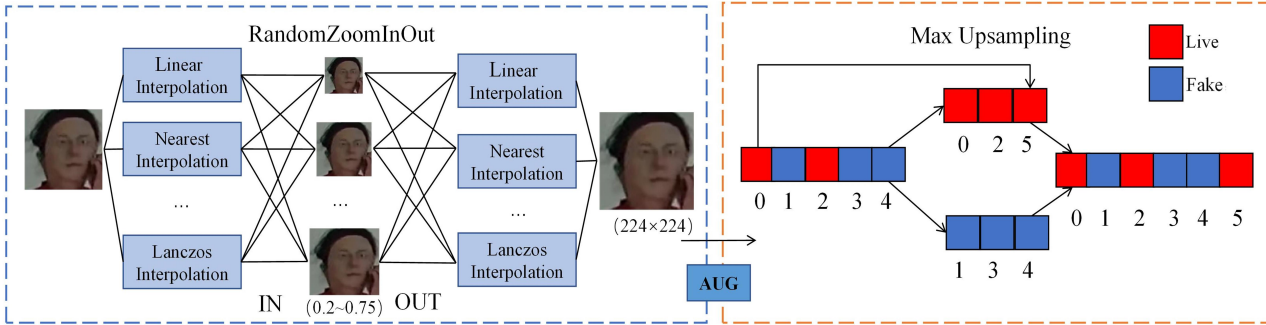


Figure 4. Pre-processing strategies used by the team of hexianhua. RandomZoomInOut data enhancement method randomly picks an interpolation method to shrink the image. The Max Upsampling method mitigates the long tail problem by complementing fewer categories. The AUG represents the various methods.

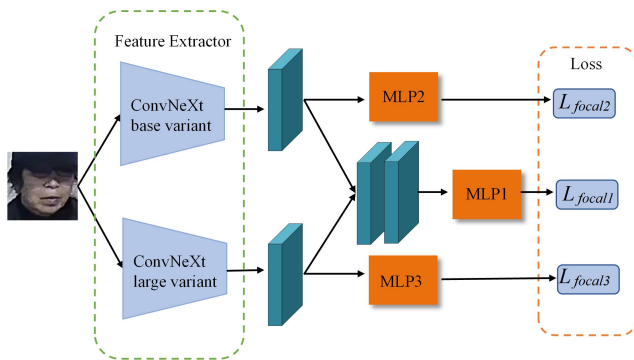


Figure 5. The framework of team OPDAI. The two branches extract features with different variants of ConvNextv2 respectively, and subsequently, the two features are concatenated as fused features. Finally, the three features are sent to separate MLP modules and supervised with three focal losses.

to obtain the appropriate parameters after several iterations helps to quickly mine the hard samples in the dataset that are difficult to discriminate. (2) Fine-tuning the whole network framework after pre-training enables the encoder to retain the ability to classify easy samples. In addition, they experimentally demonstrate that the AdamW optimizer and TTA (test time augmentation) can achieve superior performance when dealing with complex and variable images.

3.6. SeaRecluse

Since the challenge datasets contain multiple sizes of images, the team of SeaRecluse tries to find a uniform size of the input images that can maximize the performance of the algorithm. Specifically, they trained multiple models simultaneously, and during the training of each different model, they scaled the data in advance to a fixed size such as 112×112 , 128×128 , 168×168 , 192×192 , 224×224 . Then, they tested the trained model on the validation set and recorded the score for each test image. During the statis-

tical process, they found that the model trained on data of size 224×224 had the best generalization performance for the images in the surveillance scenario.

In general, they used ConvNeXt_base [39] as the backbone network of their method and used the techniques of mix up [52] and label smoothing [44] to improve the generalization of the network to complex data distributions. For the training strategy, they choose SoftTargetCrossEntropy as the loss function and use the CyclicLRd scheduler to let the learning rate cycle between $1e-5$ and 0.002 . It is worth mentioning the data enhancement scheme considering that image flipping was used in the training phase. During the model inference, they weighted the prediction scores of the original and flipped images as the final prediction result.

3.7. XiangR

In order to explore the association between model size and algorithm performance, the XiangR team at Xiamen University trained ConvNext variants with different parameter scales (ConvNext-T, ConvNext-B, ConvNext-L, ConvNext-XL) under the same experimental setup. They found that the larger model performed better on the dataset of this challenge, but constrained by hardware resources and time constraints, they chose ConvNext-B as the solution backbone network.

Before starting the training, they used various pre-processing methods such as Random Resized Crop, Horizontal Flip, and Random Affine to increase the diversity of the data. Then, he used Mix-up, Transfer Learning, Label Smoothing, and other tricks to improve the model's ability to generalize data with the complex distribution. For the training strategy, they chose the AdamW optimizer, which is initialized with a learning rate of 0.01 and a weight decay factor of $1E-2$. In addition, a cosine annealing learning rate scheduler is used, which gradually decreases the learning rate in a cosine pattern. This scheduler has a maximum number of epochs of 100 and a minimum learning rate of 0 . In other words, the AdamW optimizer with weight decay

helps prevent over-fitting, while the cosine annealing scheduler gradually decreases the learning rate to accelerate the model optimization. Finally, the cross-entropy loss is used to supervise the training process of the model.

3.8. cheniyifan

Considering that various patches of the face contain deceptive features, the team of cheniyifan proposes a high-resolution face patch pipeline based on the original image cropped out, as shown in Fig. 6. The networks on these different branches can focus on learning deceptive traces on a particular face patch, and these traces are used as additional information to assist in the classification of the whole model. In the ready state, the original image is resized to a size of 256×256 . The bounding box used to crop individual patches is predefined to be no more than 256×256 . The face boundaries are then used to crop four face patches: face, eyes, nose, and chin. It is worth stating that the original image is also used as a branch input. Then, each patch was resized to 224×224 and input to each network branch after mix up [52]. As shown in Fig. 6, the team of cheniyifan builds a multi-branch network, including branches for the original image, face, eyes, nose, and chin, each branch has a separate encoder to learn specific features, and the five encoders are efficientnetv2_m, efficientnetv2_b2, efficientnetv2_b1, efficientnetv2_b0, efficientnetv2_b0. Next, they input these five features into the attention module and supervise them with binary cross-entropy (BCE) loss. Finally, these five features are concatenated to obtain the fused deception information and perform classification. The final losses are as follows:

$$\mathcal{L}_{total} = 3 * \mathcal{L}_{ori} + \mathcal{L}_{face} + 0.5 * \mathcal{L}_{eyes} + 0.5 * \mathcal{L}_{nose} + 0.5 * \mathcal{L}_{chain} + 3 * \mathcal{L}_{concat} \quad (4)$$

where all the losses are BCE losses. Since patches such as the nose, eyes, and chin contain little deceptive information, they reduce the weights of these loss terms to 0.5 times to reduce the confidence of the whole network in this auxiliary information. For the training strategy, the cosine annealing schedule is used to decay the learning rate, where the initial learning rate is set to 0.0001 and the learning cycle decay is set to 0.5.

3.9. ioNetworks

The team of ioNetworks tries to use feature maps from multiple scales to obtain spoof traces. They use Deep-PixBis [13] as the basic framework, DenseNet as the backbone to reuse features from the previous layers, and fuse feature maps from multiple scales.

They utilize BCE to supervise the output feature maps and encourage the network to learn local information about the image patches. As shown in Fig. 7, to improve the robustness of the model to noise in the challenge dataset, in-

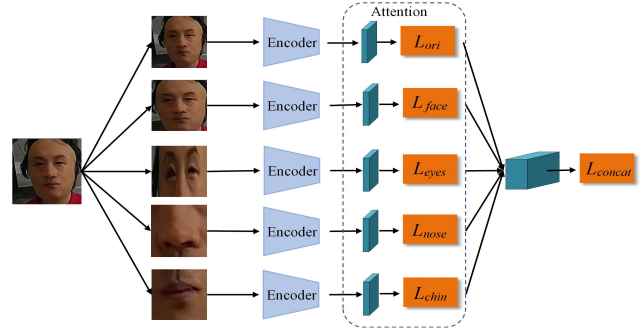


Figure 6. The pipeline of the team cheniyifan. The original image and four face patches cropped with prior knowledge are sent to each of the five encoders for feature extraction. The five features are then sent to the attention module and are concatenated as fused features. Finally, to supervise the classification of these six features, six focal losses are employed.

spired by advanced face recognition research [7], they also employ sub-center angular margin loss [16] to supervise the binary output. This loss encourages sub-classes containing clean samples to dominate the optimization of the model thus improving the robustness of the model. The overall loss function is as follows:

$$\mathcal{L}_{total} = 0.5 * \mathcal{L}_{BCE} + 0.5 * \mathcal{L}_{Ang} \quad (5)$$

where the second term of this formula is actually a variant of binary cross-entropy loss, which is calculated as:

$$\mathcal{L}_{Ang} = -1/n \sum_i (t[i] * \cos(\theta + 0.5) + (1 - t[i]) * \log(1 - \cos(\theta))) \quad (6)$$

where $t[i]$ represents the label of the i -th sample and θ represents the angle between the features learned by the backbone network and the weights.

For the training strategy, they used AdamW as an optimizer, with a learning rate of 0.0001 and a weight decay factor of 0.0005. To adjust the learning rate during training, they utilized a scheduler that decreased the learning rate for each parameter group by a factor of 0.8 every 20 epochs. Finally, they implemented early stopping with patience of 50 epochs to halt model training at a better convergence time.

4. Challenge Results

4.1. Challenge Results Report

We adopted four metrics to evaluate the performance of the solutions, which are APCER, NPCER, ACER, and AUC respectively. Please note that although we report performance on various evaluation metrics, ACER is the leading metric that determines the final ranking of the participating

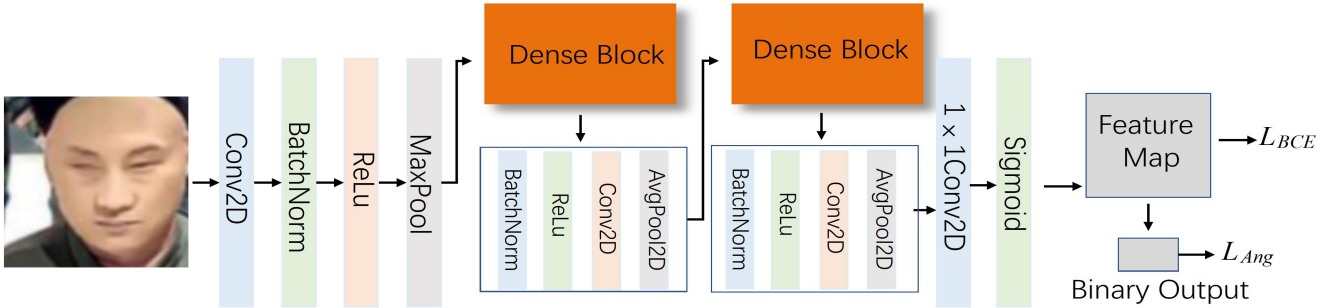


Figure 7. The framework of the team ioNetworks. A DeepPixBis-based framework is used to extract features. Binary cross-entropy loss and sub-center angular margin loss are used simultaneously to optimize the network.

teams. As shown in Tab. 3, which lists the results and ranking of the top 9 teams. We can draw the following three conclusions: (1) The first-place team got the best results on indicators ACER, APCER, AUC, and the second-place team got the best results on BPCER. (2) Eight of the top nine teams are from the industry, indicating that the task of attack detection in surveillance scenarios is important in real-world applications. (3) The top nine teams’ ACER scores were distributed between 4% and 12% and there was some disparity in the results between teams, which shows that face attack detection under surveillance is challenging and proves the value of SuHiFiMask’s in-depth research.

Table 3. Team and results are listed in the final ranking of this challenge.

R.	Team	ACER(%)	APCER(%)	BPCER(%)	AUC(%)
1	MateoH	4.73	5.07	4.38	98.38
2	CTEL_AI	5.56	9.20	1.91	98.21
3	horsego	6.22	8.17	4.26	96.97
4	hexianhua	7.08	11.21	2.94	97.83
5	OPDAI	7.16	9.18	5.13	97.38
6	SeaRecluse	9.89	15.92	3.86	96.02
7	XiangR	9.96	11.35	8.57	95.80
8	Chenyifan	11.01	14.12	7.89	94.39
9	ioNetworks	12.00	15.47	8.53	95.13

4.2. Competition summary and Future Work

Through the challenge, we summarize the effective ideas for surveillance face presentation attack detection: (1) For backbone networks, the performance is superior for networks with a large number of parameters, which proves that large models play an important role in complex applications in real life. (2) At the data level, using pre-processing methods to increase the diversity of data can effectively prevent model over-fitting, and the method of balancing the num-

ber of categories can improve the stability of the algorithm. (3) Multi-branching based feature learning method has been shown by several teams to be an effective approach for FAS tasks. In the following work, we further improve the performance from the following aspects: (1) We will explore more efficient super-resolution [3, 46, 47] methods to recover detailed features of low-quality face images. (2) We will generate [15] multiple images with similar data distribution as SuHiFiMask to guide the model to learn more general features. (3) We will explore the interpretability and validity of the large model for FAS tasks.

5. Conclusion

We organized the *Surveillance Face Presentation Attack Detection Challenge at CVPR2023* based on the SuHiFiMask dataset and running on the CodaLab platform. 180 teams registered for the competition and 37 teams made it to the final stage. In the final stage of the competition, a total of nine teams submitted their code and participated in the ranking of results after verification by the organizers. We first introduce the associated dataset, the competition protocol, and the evaluation metrics. Then, we review the solutions of the participating ranked teams and report the results of the final phase. Finally, we summarize the conclusions related to the challenges and point out effective methods for attack detection for remote monitoring scenarios.

Acknowledgments

This work was supported by the National Key Research and Development Plan under Grant 2021YFF0602103, the External cooperation key project of Chinese Academy Sciences 173211KYSB20200002, the Chinese National Natural Science Foundation Projects 61876179, 61961160704, and 62276254, the Science and Technology Development Fund of Macau Project (No.~0010/2019/AFJ, 0008/2019/A1, 0025/2019/AKP, 0019/2018/ASC), and the InnoHK program.

References

- [1] Zinelabidine Boulkenafet, Jukka Komulainen, Zahid Akhtar, Azeddine Benlamoudi, Djamel Samai, Salah Eddine Bekhouche, Abdelkrim Ouafi, Fadi Dornaika, Abdelmalik Taleb-Ahmed, Le Qin, et al. A competition on generalized software-based face presentation attack detection in mobile scenarios. In *2017 IEEE International Joint Conference on Biometrics (IJCB)*, pages 688–696. IEEE, 2017. 2
- [2] Zinelabidine Boulkenafet, Jukka Komulainen, Lei Li, Xiaoyi Feng, and Abdenour Hadid. 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)*, pages 612–618. IEEE, 2017. 1
- [3] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Basicvsr: The search for essential components in video super-resolution and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4947–4956, 2021. 8
- [4] Xiangning Chen, Chen Liang, Da Huang, Esteban Real, Kaiyuan Wang, Yao Liu, Hieu Pham, Xuanyi Dong, Thang Luong, Cho-Jui Hsieh, et al. Symbolic discovery of optimization algorithms. *arXiv preprint arXiv:2302.06675*, 2023. 5
- [5] Zhihong Chen, Taiping Yao, Kekai Sheng, Shouhong Ding, Ying Tai, Jilin Li, Feiyue Huang, and Xinyu Jin. Generalizable representation learning for mixture domain face anti-spoofing. *arXiv preprint arXiv:2105.02453*, 2021. 1
- [6] Ivana Chingovska, André Anjos, and Sébastien Marcel. On the effectiveness of local binary patterns in face anti-spoofing. In *2012 BIOSIG-proceedings of the international conference of biometrics special interest group (BIOSIG)*, pages 1–7. IEEE, 2012. 1
- [7] Jiankang Deng, Jia Guo, Tongliang Liu, Mingming Gong, and Stefanos Zafeiriou. Sub-center arcfacel: Boosting face recognition by large-scale noisy web faces. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 741–757. Springer, 2020. 7
- [8] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotzia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5203–5212, 2020. 2
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 4
- [10] Nesli Erdogmus and Sébastien Marcel. Spoofing 2d face recognition systems with 3d masks. In *2013 International Conference of the BIOSIG Special Interest Group (BIOSIG)*, pages 1–8. IEEE, 2013. 1
- [11] Hao Fang, Ajian Liu, Jun Wan, Sergio Escalera, Chenxu Zhao, Xu Zhang, Stan Z Li, and Zhen Lei. Surveillance face anti-spoofing. *arXiv preprint arXiv:2301.00975*, 2023. 2
- [12] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 4
- [13] Anjith George and Sébastien Marcel. Deep pixel-wise binary supervision for face presentation attack detection. In *2019 International Conference on Biometrics (ICB)*, pages 1–8. IEEE, 2019. 1, 7
- [14] Anjith George, Zohreh Mostaani, David Geissenbuhler, Olegs Nikisins, André Anjos, and Sébastien Marcel. Biometric face presentation attack detection with multi-channel convolutional neural network. *IEEE Transactions on Information Forensics and Security*, 15:42–55, 2019. 1
- [15] Jingwen He, Wu Shi, Kai Chen, Lean Fu, and Chao Dong. Gcfsr: a generative and controllable face super resolution method without facial and gan priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1889–1898, 2022. 8
- [16] Md Sourave Hossain, Labiba Rupty, Koushik Roy, Md Hasan, Shirshajit Sengupta, and Nabeel Mohammed. A-deeppixbis: Attentional angular margin for face anti-spoofing. In *2020 Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–8. IEEE, 2020. 7
- [17] Shan Jia, Guodong Guo, and Zhengquan Xu. A survey on 3d mask presentation attack detection and countermeasures. *Pattern recognition*, 98:107032, 2020. 1
- [18] Shan Jia, Xin Li, Chuanbo Hu, Guodong Guo, and Zhengquan Xu. 3d face anti-spoofing with factorized bilinear coding. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10):4031–4045, 2020. 1
- [19] Minchul Kim, Anil K. Jain, and Xiaoming Liu. Adafacel: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18750–18759, June 2022. 1
- [20] Pei Li, Loreto Prieto, Domingo Mery, and Patrick J. Flynn. On low-resolution face recognition in the wild: Comparisons and new techniques. *IEEE Transactions on Information Forensics and Security*, 14(8):2000–2012, 2019. 1
- [21] Xuan Li, Jun Wan, Yi Jin, Ajian Liu, Guodong Guo, and Stan Z Li. 3dpc-net: 3d point cloud network for face anti-spoofing. In *2020 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–8. IEEE, 2020. 1
- [22] Yanyu Li, Ju Hu, Yang Wen, Georgios Evangelidis, Kamyar Salahi, Yanzhi Wang, Sergey Tulyakov, and Jian Ren. Rethinking vision transformers for mobilenet size and speed. *arXiv preprint arXiv:2212.08059*, 2022. 5
- [23] Bofan Lin, Xiaobai Li, Zitong Yu, and Guoying Zhao. Face liveness detection by rppg features and contextual patch-based cnn. In *Proceedings of the 2019 3rd international conference on biometric engineering and applications*, pages 61–68, 2019. 1
- [24] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 5
- [25] Ajian Liu, Xuan Li, Jun Wan, Yanyan Liang, Sergio Escalera, Hugo Jair Escalante, Meysam Madadi, Yi Jin,

- Zhuoyuan Wu, Xiaogang Yu, et al. Cross-ethnicity face anti-spoofing recognition challenge: A review. *IET Biometrics*, 10(1):24–43, 2021. 2
- [26] Ajian Liu and Yanyan Liang. Ma-vit: Modality-agnostic vision transformers for face anti-spoofing. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 1180–1186. International Joint Conferences on Artificial Intelligence Organization, 2022. 1
- [27] Ajian Liu, Zichang Tan, Jun Wan, Sergio Escalera, Guodong Guo, and Stan Z Li. Casia-surf cefa: A benchmark for multi-modal cross-ethnicity face anti-spoofing. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1179–1187, 2021. 1
- [28] Ajian Liu, Zichang Tan, Jun Wan, Yanyan Liang, Zhen Lei, Guodong Guo, and Stan Z Li. Face anti-spoofing via adversarial cross-modality translation. *IEEE Transactions on Information Forensics and Security*, 16:2759–2772, 2021. 1
- [29] Ajian Liu, Jun Wan, Sergio Escalera, Hugo Jair Escalante, Zichang Tan, Qi Yuan, Kai Wang, Chi Lin, Guodong Guo, Isabelle Guyon, et al. Multi-modal face anti-spoofing attack detection challenge at cvpr2019. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 2
- [30] Ajian Liu, Jun Wan, Ning Jiang, Hongbin Wang, and Yanyan Liang. Disentangling facial pose and appearance information for face anti-spoofing. In *2022 26th International Conference on Pattern Recognition (ICPR)*, pages 4537–4543. IEEE, 2022. 1
- [31] Ajian Liu, Chenxu Zhao, Zitong Yu, Anyang Su, Xing Liu, Zijian Kong, Jun Wan, Sergio Escalera, Hugo Jair Escalante, Zhen Lei, et al. 3d high-fidelity mask face presentation attack detection challenge. In *ICCV Workshop*, pages 814–823, 2021. 2
- [32] Ajian Liu, Chenxu Zhao, Zitong Yu, Jun Wan, Anyang Su, Xing Liu, Zichang Tan, Sergio Escalera, Junliang Xing, Yanyan Liang, et al. Contrastive context-aware learning for 3d high-fidelity mask face presentation attack detection. *IEEE Transactions on Information Forensics and Security*, 17:2497–2507, 2022. 1
- [33] Siqi Liu, Pong C Yuen, Shengping Zhang, and Guoying Zhao. 3d mask face anti-spoofing with remote photoplethysmography. In *European Conference on Computer Vision*, pages 85–100. Springer, 2016. 1
- [34] Si-Qi Liu, Xiangyuan Lan, and Pong C Yuen. Remote photoplethysmography correspondence feature for 3d mask face presentation attack detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 558–573, 2018. 1
- [35] Si-Qi Liu, Xiangyuan Lan, and Pong C Yuen. Multi-channel remote photoplethysmography correspondence feature for 3d mask face presentation attack detection. *IEEE Transactions on Information Forensics and Security*, 16:2683–2696, 2021. 1
- [36] Yaojie Liu, Amin Jourabloo, and Xiaoming Liu. Learning deep models for face anti-spoofing: Binary or auxiliary supervision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 389–398, 2018. 1
- [37] Yaojie Liu, Joel Stehouwer, and Xiaoming Liu. On disentangling spoof trace for generic face anti-spoofing. In *European Conference on Computer Vision*, pages 406–422. Springer, 2020. 1
- [38] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 5
- [39] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. arxiv e-prints (2022). *arXiv preprint arXiv:2201.03545*, 2022. 6
- [40] Ishan Manjani, Snigdha Tariyal, Mayank Vatsa, Richa Singh, and Angshul Majumdar. Detecting silicone mask-based presentation attack via deep dictionary learning. *IEEE Transactions on Information Forensics and Security*, 12(7):1713–1723, 2017. 1
- [41] Erdogmus Nesli and Sébastien Marcel. Spoofing in 2d face recognition with 3d masks and anti-spoofing with kinect. In *IEEE 6th International Conference on Biometrics: Theory, Applications and Systems (BTAS'13)*, pages 1–8, 2013. 1
- [42] Rui Shao, Xiangyuan Lan, Jiawei Li, and Pong C Yuen. Multi-adversarial discriminative deep domain generalization for face presentation attack detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10023–10031, 2019. 1
- [43] 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
- [44] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 6
- [45] Philipp Terhorst, Jan Niklas Kolf, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Ser-fiq: Unsupervised estimation of face image quality based on stochastic embedding robustness. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5651–5660, 2020. 2
- [46] Xintao Wang, Yu Li, Honglun Zhang, and Ying Shan. Towards real-world blind face restoration with generative facial prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9168–9178, 2021. 8
- [47] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, pages 0–0, 2018. 8
- [48] Sanghyun Woo, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. Convnex v2: Co-designing and scaling convnets with masked autoencoders. *arXiv preprint arXiv:2301.00808*, 2023. 5

- [49] Bowen Yang, Jing Zhang, Zhenfei Yin, and Jing Shao. Few-shot domain expansion for face anti-spoofing. *arXiv preprint arXiv:2106.14162*, 2021. [1](#)
- [50] Zitong Yu, Jun Wan, Yunxiao Qin, Xiaobai Li, Stan Z Li, and Guoying Zhao. Nas-fas: Static-dynamic central difference network search for face anti-spoofing. *IEEE transactions on pattern analysis and machine intelligence*, 43(9):3005–3023, 2020. [1](#)
- [51] 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](#)
- [52] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. [6](#), [7](#)
- [53] Ke-Yue Zhang, Taiping Yao, Jian Zhang, Ying Tai, Shouhong Ding, Jilin Li, Feiyue Huang, Haichuan Song, and Lizhuang Ma. Face anti-spoofing via disentangled representation learning. In *European Conference on Computer Vision*, pages 641–657. Springer, 2020. [1](#)
- [54] Shifeng Zhang, Ajian Liu, Jun Wan, Yanyan Liang, Guodong Guo, Sergio Escalera, Hugo Jair Escalante, and Stan Z Li. Casia-surf: A large-scale multi-modal benchmark for face anti-spoofing. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2(2):182–193, 2020. [1](#)
- [55] Shifeng Zhang, Xiaobo Wang, Ajian Liu, Chenxu Zhao, Jun Wan, Sergio Escalera, Hailin Shi, Zezheng Wang, and Stan Z Li. A dataset and benchmark for large-scale multi-modal face anti-spoofing. In *CVPR*, 2019. [1](#)
- [56] Zhiwei Zhang, Junjie Yan, Sifei Liu, Zhen Lei, Dong Yi, and Stan Z Li. A face antispoofing database with diverse attacks. In *2012 5th IAPR international conference on Biometrics (ICB)*, pages 26–31. IEEE, 2012. [1](#)
- [57] Yaoyao Zhong, Weihong Deng, Jiani Hu, Dongyue Zhao, Xian Li, and Dongchao Wen. Sface: Sigmoid-constrained hypersphere loss for robust face recognition. *IEEE Transactions on Image Processing*, 30:2587–2598, 2021. [1](#)