



# RADAP: A Robust and Adaptive Defense Against Diverse Adversarial Patches on face recognition

Xiaoliang Liu<sup>a,b</sup>, Furao Shen<sup>a,c,\*</sup>, Jian Zhao<sup>d</sup>, Changhai Nie<sup>a,b</sup>

<sup>a</sup> National Key Laboratory for Novel Software Technology, Nanjing University, China

<sup>b</sup> Department of Computer Science and Technology, Nanjing University, China

<sup>c</sup> School of Artificial Intelligence, Nanjing University, China

<sup>d</sup> School of Electronic Science and Engineering, Nanjing University, China

## ARTICLE INFO

### Keywords:

Face recognition  
Adversarial patches  
Defense mechanism  
Deep learning  
Robustness

## ABSTRACT

Face recognition (FR) systems powered by deep learning have become widely used in various applications. However, they are vulnerable to adversarial attacks, especially those based on local adversarial patches that can be physically applied to real-world objects. In this paper, we propose RADAP, a robust and adaptive defense mechanism against diverse adversarial patches in both closed-set and open-set FR systems. RADAP employs innovative techniques, such as F-Cutout and F-patch, which use Fourier space sampling masks to improve the occlusion robustness of the FR model and the performance of the patch segmenter. Moreover, we introduce an edge-aware binary cross-entropy (EBCE) loss function to enhance the accuracy of patch detection. We also present the split and fill (SAF) strategy, which is designed to counter the vulnerability of the patch segmenter to complete white-box adaptive attacks. We conduct comprehensive experiments to validate the effectiveness of RADAP, which shows significant improvements in defense performance against various adversarial patches, while maintaining clean accuracy higher than that of the undefended Vanilla model.

## 1. Introduction

Face recognition (FR) technology is used extensively in our everyday lives, from unlocking our phones to verifying payments and identifying individuals. With the advancement of deep learning, FR has achieved remarkable progress in recent years. However, deep neural networks are also vulnerable to adversarial attacks, which can fool FR systems by adding imperceptible perturbations to the input images [1, 2]. Among various types of adversarial attacks, local-based adversarial patches are especially threatening, as they can be printed and attached to physical objects in the real world, such as glasses [3], hats [4], or masks [5,6]. These patches may lead to misclassification of the individual who wears them or exclusion of other people in the scene by the FR system. The physical realizability of adversarial patches poses a greater threat to real-world FR systems than the global perturbation-based adversarial examples. Therefore, in this paper, we focus on how to defend against adversarial patches on FR.

FR systems can be categorized as either closed-set or open-set systems. Closed-set systems aim to identify individuals from a predefined set of known subjects, while open-set systems are designed to accommodate recognition of unknown individuals not present in the predefined set. Closed-set FR can be compared to a problem of image

classification in which the goal is to categorize a face within specific and predefined categories. In contrast, open-set FR is akin to a distance metric problem, which involves the measurement of similarity between the query face and the known subjects.

While many adversarial patch defense methods primarily focus on image classification tasks, such as Interval Bound Propagation (IBP) [7], Clipped BagNet (CBN) [8], Defense against Occlusion Attacks (DOA) [9], Local Gradients Smoothing (LGS) [10], PatchGuard (PG) [11], PatchCleanser (PC) [12], it is crucial to recognize that in the real world, FR systems often operate as open-set systems. This is because they must contend with the possibility of encountering unknown individuals beyond their training data when deployed in uncontrolled environments.

**As a result, the first challenge in developing effective defense strategies is addressing both patch attacks in closed-set FR systems and open-set FR systems.** To address this challenge, we utilize a patch segmenter capable of detecting adversarial patches and subsequently concealing them using masking techniques. This process introduces four main challenges.

**Firstly, there is a need to acquire suitable training data for the patch segmenter and ensure its capability to detect adversarial**

\* Corresponding author at: School of Artificial Intelligence, Nanjing University, China.

E-mail addresses: [xiaoliang\\_liu@smail.nju.edu.cn](mailto:xiaoliang_liu@smail.nju.edu.cn) (X. Liu), [frshen@nju.edu.cn](mailto:frshen@nju.edu.cn) (F. Shen), [jianzhao@nju.edu.cn](mailto:jianzhao@nju.edu.cn) (J. Zhao), [changhainie@nju.edu.cn](mailto:changhainie@nju.edu.cn) (C. Nie).

**patches of various shapes.** To tackle this issue, we present F-patch, a method that generates patches with diverse shapes through random sampling in Fourier space. This approach equips our trained patch segmenter to recognize a broad spectrum of adversarial patches.

**Secondly, we must enhance the patch segmenter’s ability to detect patch boundaries.** To further refine the accuracy of patch detection, we propose an edge-aware binary cross-entropy (EBCE) loss function aimed at improving edge perception.

**Thirdly, as we remove adversarial patches by applying masks, it becomes imperative to fortify the occlusion robustness of the FR model.** Leveraging the well-established Cutout data augmentation technique, we introduce a mask created through random Fourier space sampling, which we term FCutout. This augmentation method significantly strengthens the model’s capacity to handle occlusions.

**Lastly, given that our patch segmenter employs deep neural networks, we must address the vulnerability to complete white-box adaptive attacks after the patch segmenter model is leaked.** To mitigate this issue, we propose the split and fill (SAF) strategy, which substantially enhances the robustness of our defense method.

Our contributions can be summarized as follows:

- We introduce RADAP, a robust and adaptive defense approach designed to counter diverse adversarial patches in both closed-set and open-set FR systems.
- To enhance the occlusion robustness of the FR model and optimize the performance of the patch segmenter, we propose two novel techniques: FCutout and F-patch. These methods employ randomly sampled masks in Fourier space. Additionally, we introduce the EBCE loss function to further enhance the patch segmenter’s performance.
- We present the SAF strategy, specifically developed to mitigate the risk of complete white-box adaptive attacks following the leakage of the segmenter model.
- Ultimately, we validate the efficacy of RADAP through a series of extensive experiments. Comparative analysis with other state-of-the-art methods demonstrates RADAP’s significant improvements in defense performance against a wide range of diverse adversarial patches. Notably, the clean accuracy of RADAP surpasses that of the undefended Vanilla model.

## 2. Related works

This section presents a review of pertinent literature concerning Adversarial Patches in the context of FR and the strategies employed for Defending Against Adversarial Patches.

### 2.1. Adversarial patches on face recognition

The primary objective of adversarial patch attacks on FR systems is to manipulate the decision-making process, potentially leading to misclassifications or unauthorized access. These attacks can be broadly categorized into two main types: evasion attacks and impersonation attacks. Evasion attacks focus on reducing the confidence in similarity measurements between identical identity pairs, resulting in misclassifications and a degradation of system accuracy. In contrast, impersonation attacks aim to enhance the feature similarity between the source images of the target individual and the various identities used by the attacker, thereby enabling unauthorized access to the FR system.

A substantial number of methods for adversarial patch attacks have been proposed, including Adv-Glasses [3], Adv-Hat [4], and Adv-Masks [6]. These methods ingeniously employ adversarial patches such as eyeglass frames, hats, and respirators, which have proven to be remarkably effective in real-world scenarios. Moreover, the transferability of adversarial patches has been significantly improved through methods like TAP [13], RSTAM [5] and EAP [14], which diversify the inputs. TAP achieves this by generating patches with a variety of

attributes and styles, allowing them to adapt to different facial features and environmental scenarios. In contrast, RSTAM leverages random similarity transformations to accommodate variations in facial poses and lighting conditions. Building on RSTAM, EAP [14] incorporates image pyramid strategies and meta-ensemble attack strategies, further enhancing the success rate of black-box attacks.

Collectively, these studies underscore the vulnerabilities inherent in facial recognition systems and emphasize the critical importance of our research on developing defenses against adversarial patches.

### 2.2. Defense against adversarial patches

Numerous researchers have dedicated their efforts to defending against adversarial patches. Among these, JPEG Compression (JPEG) [15], Digital Watermark (DW) [16], and LGS [10] emerged as early solutions. JPEG compression alleviates the impact of adversarial patches through image compression. DW identifies unnaturally dense regions in the classifier’s significance map and conceals them to prevent their influence on classification. LGS is based on the observation that patch attacks introduce concentrated high-frequency noise and proposes gradient smoothing for areas with high gradient magnitude.

However, these methods are susceptible to complete white-box adaptive attacks in the event of defense method compromise. In response to this challenge, Jiang et al. introduced IBP [7]. Similarly, the adversarial training-based [17] method DOA [9] does not suffer from this vulnerability. More recently, Xiang et al. put forth PG [11], a network characterized by small receptive fields and outlier masking. However, it requires substantial modifications to the backbone classifier. In recognition of this, Xiang et al. also introduced PC [12], a two-stage judgment-based approach applicable to any classifier. In contrast to the judgment-based PC, Segment and Complete (SAC) [18] deploys a deep neural network for adversarial patch detection and removal.

While PG and PC excel at defending against small patches, they face challenges when confronted with larger patches. To enhance defenses against adaptive attacks, SAC employs completion techniques, albeit rendering it less effective against random multiple patches. In contrast, our approach, RADAP, overcomes these limitations, providing adaptive defense against a wide array of adversarial patches.

## 3. Preliminary

### 3.1. Face recognition

In this paper, we focus on defense against adversarial patches on FR systems. We begin with a brief overview of FR systems and their tasks. As depicted in Fig. 1, FR encompasses two primary tasks: face identification and face verification. Face identification is the task of identifying a specific individual in an image by comparing it to known faces. On the other hand, face verification determines whether two given images belong to the same person. FR can be further divided into two categories: closed-set and open-set FR. In closed-set recognition, the system works with predefined identities, treating face identification as a multi-class classification problem. Neural networks used for closed-set FR are generally considered as classifiers in this context. In contrast, open-set recognition deals with unknown identities and employs metric learning techniques to measure the similarity between images in a selected feature space during face verification. Open-set face identification takes this a step further by comparing distances between the test image and each gallery image to identify the closest match. In open-set recognition, neural networks are often viewed as feature extractors.

### 3.2. Attack formulation

This paper aims to address a range of challenges in the field of FR, such as adversarial patch attacks in both closed and open set

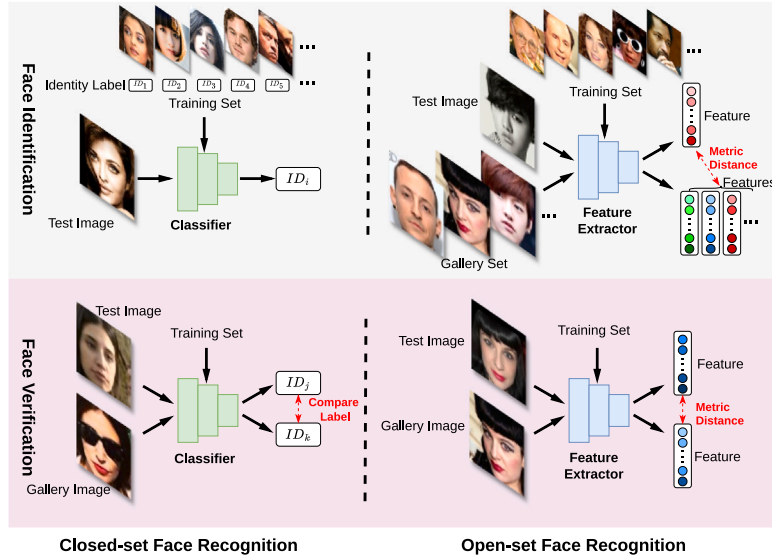


Fig. 1. Comparison of closed-set and open-set face recognition.

Table 1  
Summary of important notation.

Notation	Description
$x_a$	A attack image.
$x_s$	A attacker's source image.
$x_t$	A target image.
$y_s$	A identity label of $x_s$ .
$y_t$	A identity label of $x_t$ .
$M$	A binary mask.
$\odot$	The element-wise product operator.
$F_{FR}$	A face recognition model.
$\mathcal{L}_{CE}$	The cross-entropy loss function.
$D_{CS}$	The cosine similarity distance measure function.
$X_{ij}$	The $i, j$ th element of the two dimensional matrix $X$ .
$\otimes$	The convolutional product operator.

scenarios, as well as evasion and impersonation attacks. To delineate the adversarial patch attacks, we introduce essential concepts and variables. Specifically, we denote the attacker's source image as  $x_s$ , the target image as  $x_t$ , and the attack image as  $x_a$ . Additionally, we utilize the pre-trained FR model, represented as  $F_{FR}$ , and a binary mask, denoted as  $M$ . Furthermore, we define the cross-entropy loss function, denoted as  $\mathcal{L}_{CE}$ , as follows:

$$\mathcal{L}_{CE}(p, y) = - \sum_{i=1}^K y_i \log(p_i), \quad (1)$$

where  $K$  represents the number of classes,  $y_i$  denotes the true label for class  $i$ , and  $p_i$  indicates the predicted probability for class  $i$ . Additionally, we introduce the cosine similarity distance measure function, denoted as  $D_{CS}$ :

$$D_{CS}(v_1, v_2) = \frac{\langle v_1, v_2 \rangle}{\|v_1\| \|v_2\|}, \quad (2)$$

where  $v_1$  and  $v_2$  denote two vectors, and  $\langle \cdot, \cdot \rangle$  signifies the inner product of the vectors.

For clarity, we provide a summary of essential notations in Table 1. Additionally, we outline the process of generating adversarial patches as optimization problems in Table 2.

#### 4. Method

In this section, we introduce our defense method for FR, which is named "Robust and Adaptive Defense Against Diverse Adversarial

Table 2  
Attack formulation.

Target system	Attacker's goal	Formulation
Closed-set	Evasion	$\operatorname{argmax}_{x_a} \mathcal{L}_{CE}(F_{FR}(x_a), y_s)$
	Impersonation	$\operatorname{argmin}_{x_a} \mathcal{L}_{CE}(F_{FR}(x_a), y_t)$
Open-set	Evasion	$\operatorname{argmax}_{x_a} D_{CS}(F_{FR}(x_a), F_{FR}(x_s))$
	Impersonation	$\operatorname{argmin}_{x_a} D_{CS}(F_{FR}(x_a), F_{FR}(x_t))$

$$* \text{ s.t. } x_a \odot (1 - M) = x_s \odot (1 - M).$$

Patches" (RADAP). The core concept behind our approach revolves around identifying the locations of adversarial patches through patch segmentation and subsequently eliminating (occluding) these patches. This process necessitates solutions to four key challenges:

- Enhancing the occlusion robustness of the FR model.
- Devising a methodology for generating the adversarial patches utilized in training the patch segmenter.
- Elevating the segmentation accuracy of the patch segmenter.
- Addressing security concerns associated with the patch segmenter, especially in scenarios where it might leak information.

We outline the four fundamental steps of our method, each of which corresponds to one of the aforementioned challenges:

- **Step 1:** We train the FR model using the Fourier space sampling-based Cutout (FCutout) data augmentation method to bolster occlusion robustness (see Section 4.2).
- **Step 2:** Leveraging the FR model trained in Step 1, we generate Fourier space sampling-based adversarial patch (F-patch) images, which will be used to train the patch segmenter (see Section 4.3).
- **Step 3:** We train a patch segmenter model with the F-patch images obtained in Step 2, employing the edge-aware binary cross-entropy (EBCE) loss function, which can improve the segmentation accuracy. Using the trained patch segmenter, we can generate an initial patch mask for adversarial patch segmentation (see Section 4.4).
- **Step 4:** To further fortify the security of our approach and mitigate the risk of information leakage from the patch segmenter

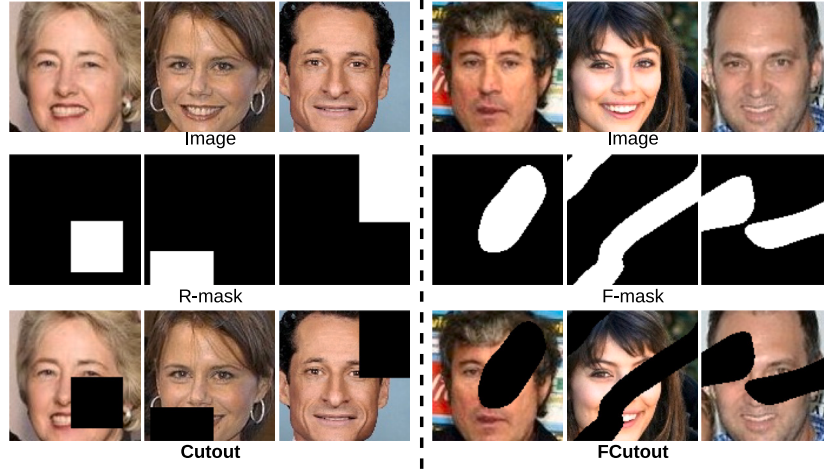


Fig. 2. Comparison of Cutout and FCutout.

model, we propose the split and fill (SAF) strategy to manipulate the initial patch masks (see Section 4.5).

FCutout, F-patch, EBCE, and SAF are proposed to address the four key challenges mentioned above. To ensure a comprehensive understanding of our methodology, we introduce the Fourier space sampling-based mask (F-mask) before delving into the four key steps (see Section 4.1).

#### 4.1. Fourier space sampling-based mask

We initiate by introducing the F-mask generation process, a technique derived from FMix [19]. This process commences with the sampling of a low-frequency grayscale mask from the Fourier space. Subsequently, the gray-scale mask undergoes a transformation into a binary form through a thresholding mechanism. To elaborate further, let us consider  $\mathbf{Z}$  as a complex random variable defined within the domain  $\mathbb{C}^{H \times W}$ . Moreover, both  $\mathcal{R}(\mathbf{Z})$  and  $\mathcal{J}(\mathbf{Z})$  follow a normal distribution  $\mathcal{N}(\mathbf{0}, \mathbf{I}_{H \times W})$ , where  $\mathcal{R}$  and  $\mathcal{J}$  denote the extraction of real and imaginary parts, respectively.

Following this, we apply a low-pass filter, denoted as  $\mathcal{F}_{LPF}$ , to  $\mathbf{Z}$ . This filter serves to attenuate the high-frequency components of  $\mathbf{Z}$  using a decay power parameter  $\delta$ . Mathematically, this attenuation can be expressed as:

$$\mathcal{F}_{LPF}(\mathbf{Z}, \delta)_{ij} = \frac{\mathbf{Z}_{ij}}{\mathcal{F}_{Freq}(H, W)_{ij}^\delta}, \quad (3)$$

where  $\mathcal{F}_{Freq}$  returns the discrete Fourier transform sample frequencies. By applying the inverse discrete Fourier transform  $\mathcal{F}_{IDFT}$  to  $\mathcal{F}_{LPF}(\mathbf{Z}, \delta)$ , we obtain a gray-scale image represented by:

$$\mathbf{G} = \mathcal{F}_{Norm}(\mathcal{R}(\mathcal{F}_{IDFT}(\mathcal{F}_{LPF}(\mathbf{Z}, \delta)))), \quad (4)$$

where  $\mathcal{F}_{Norm}$  is a normalization function, and  $\mathcal{R}$  returns the real part of the input.

The gray-scale mask,  $\mathbf{G}$ , can then be transformed into a binary mask,  $\mathbf{M}$ , using a thresholding function  $\mathcal{F}_T$ . Let  $\mathcal{F}_T([\lambda HW], \mathbf{G})$  denotes taking the  $[\lambda HW]$ -th largest value in  $\mathbf{G}$ , where  $[\cdot]$  denotes the rounding operation. We can obtain a binary mask represented by

$$\mathbf{M}_{ij} = \begin{cases} 1, & \text{if } \mathbf{G}_{ij} > \mathcal{F}_T([\lambda HW], \mathbf{G}), \\ 0, & \text{otherwise,} \end{cases} \quad (5)$$

$\lambda \sim \mathcal{U}(a, b), 0 \leq a < b \leq 1$ ,

where  $\mathcal{U}(a, b)$  denotes the uniform distribution of  $a$  to  $b$ .

#### 4.2. Step 1: Train the FR model

In this step, we train the FR model using FCutout. Our objective is to defend against adversarial patch attacks by removing (occluding) these patches, necessitating the use of occlusion-robust FR models. Among various augmentation methods applied to increase the reliability and robustness of deep networks [20], the Cutout technique is used in this work due to its efficiency in the occlusion robustness of deep neural network models [21]. However, the random rectangular mask (R-mask) employed by Cutout tends to be too uniform, whereas adversarial patch shapes exhibit diversity, resulting in various occlusion patch shapes. To address this limitation, we introduce the FCutout data augmentation method, which leverages the F-mask. Fig. 2 illustrates a comparison between FCutout and Cutout. Notably, the F-mask, unlike the R-mask, offers a wider variety of masks. This diversity results in improved occlusion robustness in FR models trained using FCutout.

#### 4.3. Step 2: Generate F-patch

In this step, we generate the adversarial patches required for training the patch segmenter. To ensure diversity in these adversarial patches, we also incorporate the use of the F-mask. Let  $\mathbf{M}$  denote the binary mask generated using the F-mask,  $\mathbf{x}_s$  represent the attacker's source image, where  $\mathbf{x}_s \in [0, 1]^{H \times W \times 3}$ . Furthermore,  $\mathcal{F}_{FR}$  denotes the FR model trained in Step 1. We employ the projected gradient descent (PGD) [2] to generate the attack image  $\mathbf{x}_a$  with F-patch, as described below:

$$\begin{aligned} \mathbf{x}_a^{(0)} &= \mathbf{x}_s, \\ \mathbf{x}_a^{(t+1)} &= \prod_{\mathbb{X}_a} (\mathbf{x}_a^{(t)} + \mathbf{M} \odot (\alpha \text{Sign}(\nabla_{\mathbf{x}_a^{(t)}} \mathcal{L}_{CE}(\mathcal{F}_{FR}(\mathbf{x}_a^{(t)}), y_s))), \\ \mathbf{x}_a &= \mathbf{x}_a^{(T)}, T \sim \mathcal{U}(1, 1000), \end{aligned} \quad (6)$$

where  $\alpha$  represents the perturbation step size,  $T$  is the iteration number. To further generating diversity  $\mathbf{x}_a$ ,  $T$  is randomly sampled from a uniform distribution from 1 to 1000. The  $\prod$  symbol denotes the projection function, and  $\mathbb{X}_a = \{\mathbf{x} : \|\mathbf{x} - \mathbf{x}_s\| \leq \epsilon \text{ and } \mathbf{x} \in [0, 1]^{H \times W \times 3}\}$ . Fig. 3 provides examples of F-patch.

#### 4.4. Step 3: Train the patch segmenter

In this step, we train the patch segmenter with the F-patch using the EBCE loss function. We use U-Net [22] as the network architecture for the patch segmenter. Let  $\mathcal{F}_{PS}$  denote the patch segmenter and  $\theta$  denote the parameters of the network.  $\mathbf{x}_a$  denotes an attack image with F-patch and  $\mathbf{M}$  is its ground-truth binary patch mask.  $\mathcal{L}_{EBCE}$  denotes the EBCE loss function. We treat the challenge of detecting adversarial



Fig. 3. Examples of F-patch.



(a)  $M$                       (b)  $\nabla M$

Fig. 4. An example of  $\nabla M$ .

patches as a segmentation problem, and our solution involves training a dedicated patch segmenter. This training process can be formulated as an optimization problem:

$$\min_{\theta} \mathcal{L}_{EBCE}(\mathcal{F}_{PS}(x_a; \theta), \mathbf{M}), \quad (7)$$

where  $\mathcal{F}_{PS}(x_a; \theta) \in [0, 1]^{H \times W}$  is the output probability map. To enhance the accuracy of detecting patch edges, we propose the EBCE loss function, which is formulated as follows:

$$\mathcal{L}_{EBCE}(\hat{\mathbf{M}}, \mathbf{M}) = -\frac{1}{HW} \sum_i^H \sum_j^W e^{\beta(\nabla M)_{ij}} [\mathbf{M}_{ij} \log \hat{\mathbf{M}}_{ij} + (1 - \mathbf{M}_{ij}) \log(1 - \hat{\mathbf{M}}_{ij})], \quad (8)$$

where  $\nabla M$  denotes the gradient of  $M$ . Fig. 4 shows an example of  $\nabla M$ . The horizontal and vertical Sobel convolution kernels are represented by  $S_x$  and  $S_y$ , respectively. The gradient magnitude  $\nabla M$  is computed as follows:

$$\nabla M = \sqrt{(\mathbf{M} \otimes S_x)^2 + (\mathbf{M} \otimes S_y)^2}, \quad (9)$$

$$S_x = \begin{pmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{pmatrix}, S_y = \begin{pmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{pmatrix},$$

where  $\otimes$  denotes the convolutional product operator. After training the patch segmenter  $\mathcal{F}_{PS}$ , we can generate the initial binary patch mask, denoted as  $M_{PS}$ , by binarizing the output probability map as follows:

$$M_{PS} = \mathcal{F}_{FR}(x_a) > 0.5. \quad (10)$$

Fig. 5 presents instances of both non-adaptive and adaptive attacks. In a non-adaptive attack scenario, the attacker solely relies on the FR model, and our patch segmenter efficiently identifies the adversarial patches, leading to a successful defense. However, it is essential to note that the patch segmenter  $\mathcal{F}_{PS}$  also employs a neural network architecture, making it susceptible to adversarial examples. This vulnerability becomes particularly pronounced in the context of strong white-box adaptive attacks where the attacker possesses comprehensive knowledge of the system. In this paper, we employ the BPDA adaptive attack [23], which involves approximating the non-differentiable part of the defense system with a differentiable approximation to execute

an effective adaptive attack. To achieve this, we use Sigmoid functions, denoted as  $\sigma$ , to approximate the non-differentiable binarization operation, as shown in the equation below:

$$\sigma(\mathcal{F}_{FR}(x_a)) \approx (\mathcal{F}_{FR}(x_a) > 0.5). \quad (11)$$

As demonstrated in Fig. 5, the adaptive attack scenario exposes a critical vulnerability in our defense system when relying solely on a single patch segmenter. To address this challenge effectively, we introduce the SAF strategy, which is explained in Section 4.5.

#### 4.5. Step 4: Using the SAF strategy

In this step, we introduce the SAF strategy to handle the initial patch mask. As depicted in Fig. 6, the SAF strategy serves as an effective defense against strong adaptive attacks, even when faced with scenarios involving potential patch segmenter information leakage.

The SAF strategy entails several key steps: Initially, the initial patch mask is divided into smaller subgrids. Next, the sum of values within each subgrid is calculated. Subgrids with a sum exceeding a predefined threshold ( $\tau$ ) are populated with 1 s, while those below the threshold are filled with 0 s. This process is applied iteratively to all subgrids, resulting in the final updated mask.

The adaptability of the SAF strategy to various adversarial patch shapes and quantities is evident in Fig. 6, highlighting its robustness against adaptive attacks. A comprehensive understanding of the SAF algorithm can be found in Algorithm 1. It is noteworthy that the overall time complexity of the SAF algorithm is  $\mathcal{O}(HW)$ .

## 5. Experiments

In this section, we comprehensively evaluate the performance of RADAP in defending against a variety of adversarial patches on FR systems. Additionally, we assess the robustness of RADAP against adaptive attacks, especially in the case where the patch segmenter model is leaked. The experiments conducted aim to provide a thorough understanding of the effectiveness and capabilities of RADAP in FR systems.

### 5.1. Experimental setup

**Datasets.** For closed-set FR systems, we use a subset of the VG-Face2 dataset [24]. The subset contains 100 classes, and we split the images into training, validation, and testing sets in a 7:1:2 ratio. For open-set FR systems, we use the MS-Celeb-1M dataset [25] for training and the LFW dataset [26] for testing. Face images from all datasets are aligned to a size of  $112 \times 112$ .

**Attack Patch Masks.** Fig. 7 illustrates the patch masks used in our testing. These masks fall into two categories: a predefined set containing masks such as glasses, stickers, and a respirator mask, and a second category of randomly generated masks, referred to as the Random Rectangle Mask (R-mask) and the Fourier space sampling-based mask (F-mask).

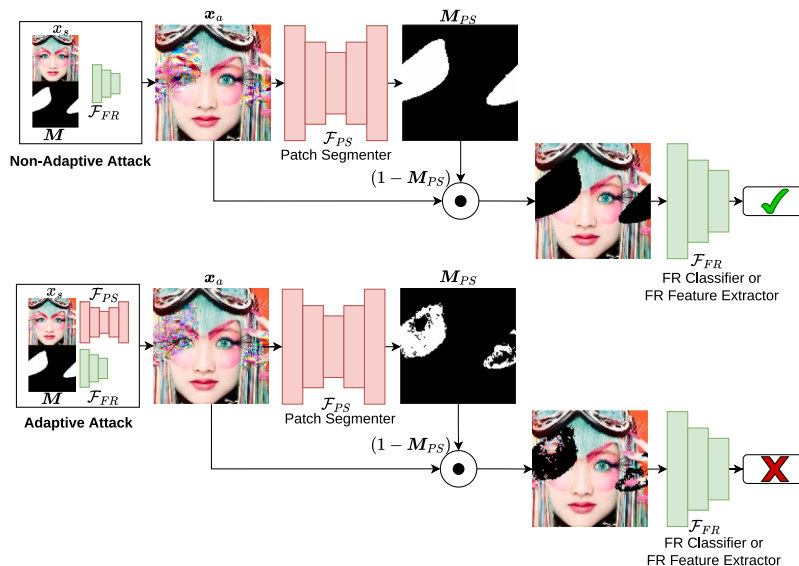


Fig. 5. Examples of non-adaptive and adaptive attacks.

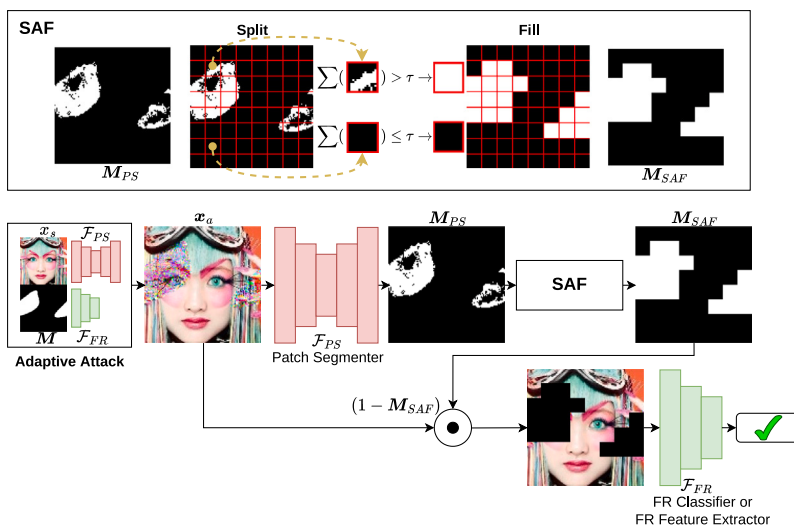


Fig. 6. Split and Fill Strategy.  $\tau$  denotes a predefined threshold.




Fig. 7. Examples of the patch masks for testing.

**Target Models.** For closed-set FR systems, we use four well-known target models, including MobileFaceNet [27], IResNet50 [31], EfficientNet-B0 [33], and Swin-T [32]. For open-set FR systems, we employ five target models, summarized in Table 3.

**Attack Methods.** In this experiment, we employ four distinct adversarial patch attack methods to evaluate the effectiveness of our defense approach. These methods consist of PGDAP [2], APGDAP [34], TAP [13], and RSTAM [5]. Specifically, PGDAP and APGDAP are the adversarial patch versions of the PGD [2] and Auto-PGD [34] attacks.

TAP and RSTAM leverage input diversity to enhance the transferability of adversarial patches, endowing them with more robust attack capabilities compared to PGDAP.

**Baseline Defense Methods.** We report the results of our method with SAF, Ours(+), and without SAF, Ours(-). Additionally, we compare our approach with five baseline adversarial patch defense methods: Defense against Occlusion Attacks (DOA) [9], JPEG Compression (JPEG) [15], Local Gradients Smoothing (LGS) [10], PatchCleanser (PC) [12], and Segment and Complete (SAC) [18].

**Algorithm 1** Split and Fill algorithm

---

**Require:** Initial patch mask  $\mathbf{M} \in [0, 1]^{H \times W}$ , with both the height and width ( $H$  and  $W$ ) being equal.

**Require:** Number of subgrids  $n$ .

- 1: Extract dimensions:  $H$  and  $W$  from  $\mathbf{M}$ .
- 2: Calculate subgrid side length:  $g_{size} = \lfloor \frac{H}{n} \rfloor$ .
- 3: Set the predefined threshold:  $\tau = g_{size}$ .
- 4: **for**  $i$  in  $[0, n)$  **do**
- 5:   **for**  $j$  in  $[0, n)$  **do**
- 6:     Compute subgrid boundaries:
- 7:      $g_h = \min((i + 1) \cdot g_{size}, H)$ ,
- 8:      $g_w = \min((j + 1) \cdot g_{size}, W)$ .
- 9:     Calculate subgrid sum:
- 10:      $g_{sum} = \sum_{x=i \cdot g_{size}}^{i \cdot g_{size} + g_h - 1} \sum_{y=j \cdot g_{size}}^{j \cdot g_{size} + g_w - 1} \mathbf{M}_{xy}$ .
- 11:     Update subgrid:
- 12:     **if**  $g_{sum} > \tau$  **then**
- 13:       Fill all elements of  $\mathbf{M}[i \cdot g_{size} : g_h, j \cdot g_{size} : g_w]$  with 1.
- 14:     **else**
- 15:       Fill all elements of  $\mathbf{M}[i \cdot g_{size} : g_h, j \cdot g_{size} : g_w]$  with 0.
- 16:     **end if**
- 17:   **end for**
- 18: **end for**
- 19: **return** updated mask  $\mathbf{M}$

---

**Table 3**  
Target models of open-set FR systems.

Target model	Neural architecture	Loss
ArcMob	MobileFaceNet [27]	ArcFace [28]
MagMob	MobileFaceNet [27]	MagFace [29]
MvMob	MobileFaceNet [27]	MV Softmax [30]
MvRes50	IResNet50 [31]	MV Softmax [30]
MvSwinT	Swin-T [32]	MV Softmax [30]

**Evaluation Metric.** For closed-set FR systems, our primary performance measure is classification accuracy (Acc). For open-set FR systems, we evaluate using True Acceptance Rate at a False Acceptance Rate of 0.001 (TAR@0.001FAR).

**Implementation Details.** During the training of the robust FR model, we set the values of  $(a, b)$  as  $(0, 1)$ . When generating the F-patch, we configure the parameters  $(a, b)$  to be  $(0.02, 0.3)$ . The patch segmenter is trained using the Stochastic Gradient Descent optimizer [35] with an initial learning rate of 0.1, weight decay of  $10^{-4}$ , momentum of 0.9, and Nesterov momentum enabled. We set the weight  $\beta$  in EBCE to 1. Moreover, we apply a cosine annealing scheduler technique [36] to adjust the learning rate schedule. For the attack, the default attack method is PGDAP, the default number of perturbation iterations is set to 100 steps, using a perturbation step size  $\alpha$  of 0.007, and the perturbation boundary  $\epsilon$  is defined as 0.3. The default value for the number of subgrids  $n$  in the SAF strategy is set to 8.

Specifically, the attack method for generating the F-patch used to train the patch segmenter exclusively employs PGDAP, with the attacker’s goal being evasion. Moreover, the target FR model is MobileFaceNet, which is trained on a closed-set dataset.

## 5.2. Comparison studies on closed-set FR systems

**Defense against Unseen Patch Masks.** Table 4 presents a comparative analysis of various defense methods against unseen patch masks in closed-set FR systems. In practice, attackers often use undisclosed patch shapes, making it crucial to establish robust defenses against a variety of shapes. Our proposed method consistently demonstrates strong defense capabilities against different patch masks, whether they are predetermined or randomly generated. Notably, it outperforms other methods, particularly in scenarios involving Glasses masks and

F-masks. For Glasses masks, “Ours(–)” outperforms “SAC” by 60.63% (91.39% vs. 30.76%), and “Ours(+)” surpasses “SAC” by 53.96% (84.72% vs. 30.76%). Additionally, our method achieves the highest performance in clean classification accuracy, indicating minimal information loss when defending against clean data.

Fig. 8 provides a comprehensive comparison of defense methods across various unseen patch masks with different parameters. Our method consistently demonstrates robust defense performance across different unseen patch masks and parameters. Notably, For R-masks, where SAC-M denotes the SAC version that specifies the number of holes. However, in reality, the number of holes employed by an attacker is often unknown to us. Our method’s adaptability eliminates the need to specify the number of holes, ensuring effective defense regardless of the attacker’s choices.

**Defense against Unseen Attacker’s Goals.** Table 4 also presents results when the attacker’s goal is impersonation. Although our patch segmenter was not specifically trained for impersonation attacks, our method exhibits superior performance in these scenarios, showcasing its adaptability.

**Defense against Unseen Target Models.** In our approach, we utilize the MobileFaceNet as the target model for generating the F-patch used in patch segmenter training. In Table 5, we present a comparative analysis of our defense method’s performance when confronted with previously unseen target models, including the transformer-based model Swin-T. Remarkably, our defense method consistently delivers strong performance across different target models.

**Defense against Unseen Attack Methods.** Table 6 presents the defense results when subjected to three attack methods stronger than PGDAP, specifically APGDAP, TAP, and RSTAM. The noteworthy performance of our method in defending against these advanced attacks underscores its remarkable robustness.

## 5.3. Ablation studies on closed-set FR systems

In this subsection, we perform ablation studies on closed-set FR systems, focusing on three key components: FCutout, F-patch, and EBCE.

**Ablation Studies with FCutout.** We start by examining the results of ablation studies with FCutout, and the findings are presented in Table 7. The table clearly indicates that FCutout outperforms both Cutout and Vanilla in terms of clean classification accuracy. This observation highlights FCutout’s superior generalization capabilities when compared to Cutout. Notably, FCutout consistently achieves higher classification accuracy than Cutout for most masks, underscoring its exceptional occlusion robustness.

**Ablation Studies with F-patch and EBCE.** Next, we delve into the results of ablation studies involving F-patch and EBCE, as detailed in Table 8. In these studies, EBCE stands out by surpassing the performance of BCE, demonstrating its superior capabilities. Additionally, when comparing F-patch to R-patch, we find that F-patch generates more diverse patches, which significantly contributes to the improved generalization capacity of our defense method. For a more visual representation of these findings, we provide visualization results in Fig. 9, clearly illustrating the effectiveness of “R-patch+EBCE” in managing patch edges compared to “R-patch+BCE”. It is worth noting that the combination of “F-patch” and “EBCE” yields the highest performance.

These ablation studies offer valuable insights into the individual contributions of FCutout, F-patch, and EBCE to the overall defense mechanism, confirming their significance in enhancing the robustness of our system.

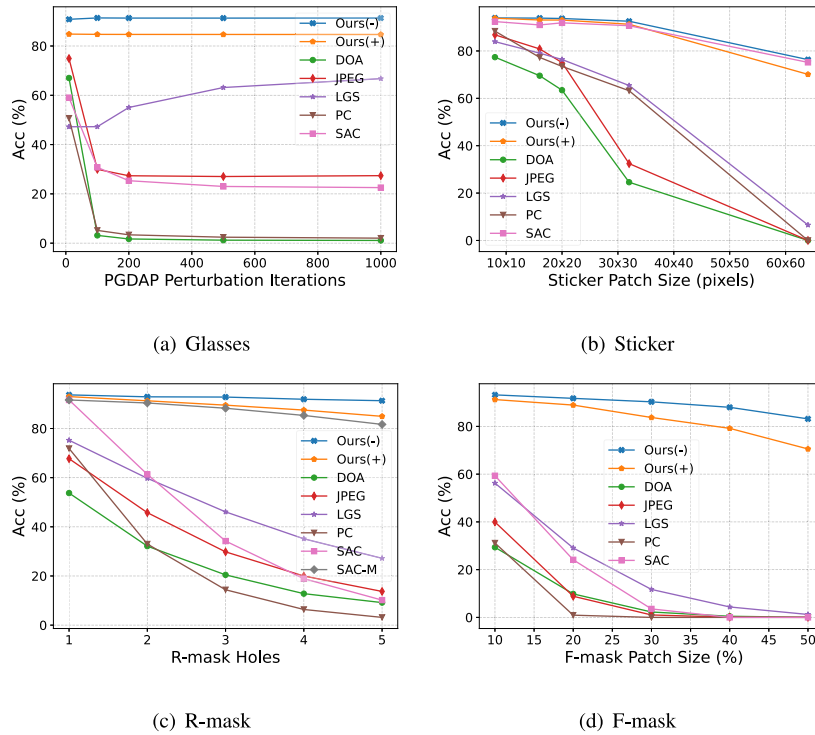
## 5.4. Generalization to open-set FR systems

In contrast to closed-set FR systems, open-set FR systems find more frequent applications, particularly in commercial systems. In our

**Table 4**

Comparison results of different methods for defending unseen patch masks (shapes) and attacker’s goals on closed-set FR systems. The best performances in each block are shown in **bold**. The evaluation metric uses Acc (%).

Attacker’s goal	Defense method	Clean	Mask				
			Glasses	Sticker	Respirator	R-mask	F-mask
Evasion	DOA [9]	80.07	3.15	24.58	1.20	30.77	12.47
	JPEG [15]	88.57	29.96	32.44	0.02	39.90	14.19
	LGS [10]	87.59	47.26	60.75	13.35	54.90	22.40
	PC [12]	92.70	5.18	63.21	0.44	49.30	12.10
	SAC [18]	92.97	30.76	90.66	79.24	88.98	21.83
	<b>Ours(-)</b>	94.00	<b>91.39</b>	<b>92.55</b>	<b>87.46</b>	<b>92.11</b>	<b>78.86</b>
	<b>Ours(+)</b>	<b>94.01</b>	84.72	91.28	83.55	89.97	69.79
Impersonation	DOA [9]	-	29.21	63.73	16.35	50.15	22.59
	JPEG [15]	-	57.00	56.74	11.43	52.15	24.39
	LGS [10]	-	68.04	77.74	43.72	72.24	39.44
	PC [12]	-	30.69	68.10	20.55	63.52	17.85
	SAC [18]	-	61.99	90.82	80.14	89.53	30.50
	<b>Ours(-)</b>	-	<b>91.42</b>	<b>92.55</b>	<b>87.46</b>	<b>92.23</b>	<b>78.79</b>
	<b>Ours(+)</b>	-	84.89	91.28	84.23	90.24	70.56

**Fig. 8.** Comparison results under different parameters for different unseen patch masks.

approach, we opted not to retrain the patch segmenter but instead transferred it directly to open-set FR systems. Table 9 provides an overview of the defense results when generalized to open-set FR systems. Once more, when confronted with unseen target models and unfamiliar patch shapes, our method exhibits superior defense performance. This observation underscores our approach’s ability to maintain robust defense performance in open-set FR systems.

### 5.5. Adaptive attacks after defense method leakage

To assess the defense method’s robustness, we conduct evaluations under adaptive attacks, considering scenarios where the attackers possess comprehensive knowledge about the FR systems due to defense method leakage. In our experiments, we employ the BPDA adaptive attack, which employs a differentiable operation to approximate the non-differentiable component of the original pipeline, enabling adaptive attacks. For “JPEG” and “LGS”, we assume the output of each defense is approximately equal to the original input. In the case of

“SAC” and “Ours”, we utilize a Sigmoid function to approximate the non-differentiable binarization operation. The results for defense against adaptive attacks are presented in Table 10. The table clearly indicates that “JPEG”, “LGS”, and “Ours(-)” exhibit weak defense against adaptive attacks, while “Ours(+)” demonstrates the highest defense performance across all patch masks.

Fig. 10 provides a comparative analysis of defense results against non-adaptive and adaptive attacks. It is clear that the “DOA” method, based on adversarial training, and the “PC” defense method, which relies on two-stage judgments, do not suffer from gradient obfuscation issues. Consequently, they exhibit consistent defense performance under both adaptive and non-adaptive attacks. “Ours(+)” also maintains similar defense performance in both scenarios, highlighting the robustness of our defense method.

Fig. 11 offers visual results of the defense masks under adaptive attacks. The “Glasses” example in Fig. 11 illustrates SAC’s subpar defense when facing patches with complex shapes, and the “F-mask” demonstrates SAC’s inadequacy when dealing with patches featuring random

**Table 5**

Comparison results of different methods for defending unseen target models on closed-set FR systems.

Target model	Defense method	Clean	Mask				
			Glasses	Sticker	Respirator	R-mask	F-mask
IResNet50 [31]	JPEG [15]	90.75	24.47	32.27	0.24	40.93	14.55
	LGS [10]	91.46	55.04	67.66	18.63	64.01	28.23
	PC [12]	94.52	2.08	73.89	0.64	56.88	14.14
	SAC [18]	95.12	24.75	93.59	87.25	92.04	24.81
	<b>Ours(-)</b>	<b>95.13</b>	<b>93.87</b>	<b>93.89</b>	<b>89.95</b>	<b>93.73</b>	<b>83.78</b>
	<b>Ours(+)</b>	94.99	89.54	92.87	86.77	92.14	74.62
EfficientNet-B0 [27]	JPEG [15]	90.80	34.80	50.36	5.29	52.38	21.99
	LGS [10]	86.16	59.82	74.89	48.23	73.30	41.40
	PC [12]	93.83	0.37	68.04	0.41	51.12	12.26
	SAC [18]	94.39	11.61	92.81	83.86	90.34	21.81
	<b>Ours(-)</b>	<b>94.48</b>	<b>92.27</b>	<b>92.96</b>	<b>88.51</b>	<b>92.48</b>	<b>81.12</b>
	<b>Ours(+)</b>	94.22	86.05	91.77	84.86	90.50	72.26
Swin-T [32]	JPEG [15]	82.79	43.54	61.53	6.44	51.11	20.63
	LGS [10]	89.89	42.95	65.51	10.15	54.55	21.34
	PC [12]	92.14	0.43	73.50	0.00	37.68	10.26
	SAC [18]	92.38	18.94	91.28	82.16	89.04	20.99
	<b>Ours(-)</b>	<b>92.41</b>	<b>91.25</b>	<b>91.51</b>	<b>87.06</b>	<b>91.07</b>	<b>80.88</b>
	<b>Ours(+)</b>	91.37	87.01	90.10	81.60	89.25	72.66

**Table 6**

Comparison results of different methods for defending unseen attack methods on closed-set FR systems.

Attack method	Defense method	Clean	Mask				
			Glasses	Sticker	Respirator	R-mask	F-mask
APGDAP [34]	DOA [9]	80.07	12.42	36.66	15.18	37.76	20.61
	JPEG [15]	88.57	38.18	47.28	8.07	46.10	19.28
	LGS [10]	87.59	71.46	79.70	52.34	74.86	42.29
	PC [12]	92.70	28.95	66.92	9.06	54.51	19.60
	SAC [18]	92.76	42.82	91.13	79.14	89.38	29.04
	<b>Ours(-)</b>	94.00	<b>91.66</b>	<b>92.90</b>	<b>88.49</b>	<b>92.36</b>	<b>79.87</b>
<b>Ours(+)</b>	<b>94.01</b>	85.69	91.91	84.51	90.73	71.12	
TAP [13]	DOA [9]	80.14	4.00	31.24	2.36	33.16	13.94
	JPEG [15]	88.43	6.99	13.27	0.00	28.59	9.31
	LGS [10]	87.59	45.84	47.72	1.88	43.22	15.88
	PC [12]	92.68	12.02	63.72	0.67	50.88	13.74
	SAC [18]	92.12	42.40	90.30	78.93	89.08	24.77
	<b>Ours(-)</b>	94.00	<b>89.96</b>	<b>92.46</b>	<b>87.40</b>	<b>92.05</b>	<b>78.13</b>
<b>Ours(+)</b>	<b>94.01</b>	84.61	91.31	82.99	90.01	69.39	
RSTAM [5]	DOA [9]	80.14	13.34	51.13	13.59	42.55	19.28
	JPEG [15]	88.43	20.22	37.50	0.05	39.96	14.96
	LGS [10]	87.59	48.96	57.22	5.14	50.77	20.47
	PC [12]	92.68	37.21	66.96	4.58	58.13	19.81
	SAC [18]	92.12	53.98	90.49	79.22	89.23	30.45
	<b>Ours(-)</b>	94.00	<b>90.90</b>	<b>92.49</b>	<b>87.39</b>	<b>92.00</b>	<b>78.37</b>
<b>Ours(+)</b>	<b>94.01</b>	84.68	91.31	83.18	90.04	69.82	

**Table 7**

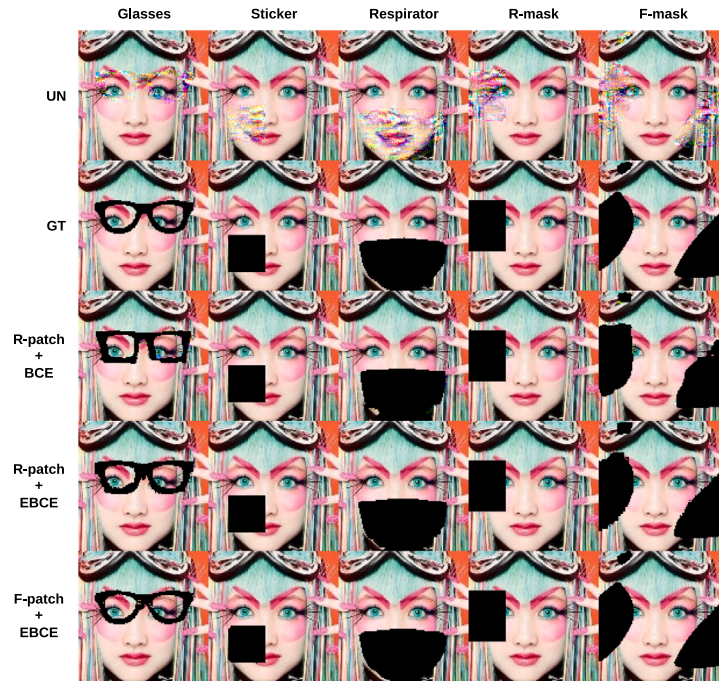
Results of ablation studies with FCutout. UN means undefended and GT is defended using ground truth mask.

Defense method		Clean	Mask				
			Glasses	Sticker	Respirator	R-mask	F-mask
UN	Vanilla	89.43	0.00	0.04	0.00	14.12	5.51
	Cutout	92.97	0.00	0.05	0.00	13.85	5.71
	FCutout	<b>94.09</b>	0.00	<b>0.32</b>	0.00	<b>17.22</b>	<b>6.40</b>
GT	Vanilla	89.43	64.23	82.96	61.95	79.27	55.32
	Cutout	92.97	88.70	90.83	83.08	90.82	71.93
	FCutout	<b>94.09</b>	<b>91.42</b>	<b>92.56</b>	<b>87.46</b>	<b>92.12</b>	<b>78.90</b>
SAC [18]	Vanilla	89.31	3.69	82.21	48.92	72.20	17.58
	Cutout	92.97	<b>30.76</b>	90.66	79.24	88.98	21.83
	FCutout	<b>94.02</b>	29.50	<b>92.49</b>	<b>83.47</b>	<b>90.17</b>	<b>24.09</b>
<b>Ours(-)</b>	Vanilla	89.43	63.81	82.97	61.94	79.29	55.06
	Cutout	92.97	88.62	90.85	83.08	90.79	71.89
	FCutout	<b>94.00</b>	<b>91.39</b>	<b>92.55</b>	<b>87.46</b>	<b>92.11</b>	<b>78.86</b>
<b>Ours(+)</b>	Vanilla	88.62	48.09	74.96	51.38	73.51	42.64
	Cutout	92.86	<b>85.14</b>	89.60	79.78	88.90	65.19
	FCutout	<b>94.01</b>	84.72	<b>91.28</b>	<b>83.55</b>	<b>89.97</b>	<b>69.79</b>

**Table 8**

Results of ablation studies with F-patch and EBCE.

Defense method	Clean	Mask				
		Glasses	Sticker	Respirator	R-mask	F-mask
R-patch + BCE	<b>94.08</b>	78.55	92.53	86.47	92.11	65.87
R-patch + EBCE	94.02	<b>90.64</b>	<b>92.53</b>	<b>87.25</b>	<b>92.12</b>	<b>77.57</b>
F-patch + BCE	93.93	91.20	92.30	87.34	92.05	78.76
F-patch + EBCE	<b>94.00</b>	<b>91.39</b>	<b>92.55</b>	<b>87.46</b>	<b>92.11</b>	<b>78.86</b>

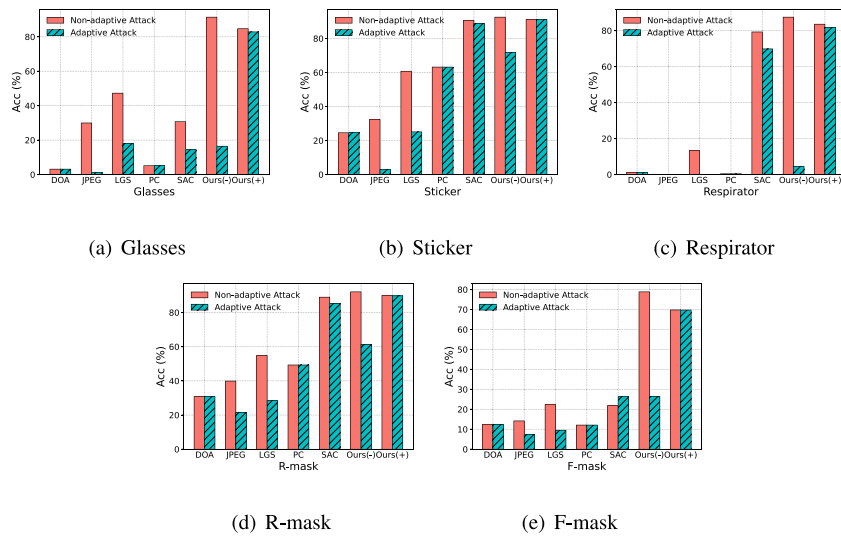
**Fig. 9.** Visualization results of ablation studies with F-patch and EBCE.**Table 9**

Evaluation of defense results in the context of generalization to open-set FR systems, utilizing the TAR@0.001FAR (%) evaluation metric.

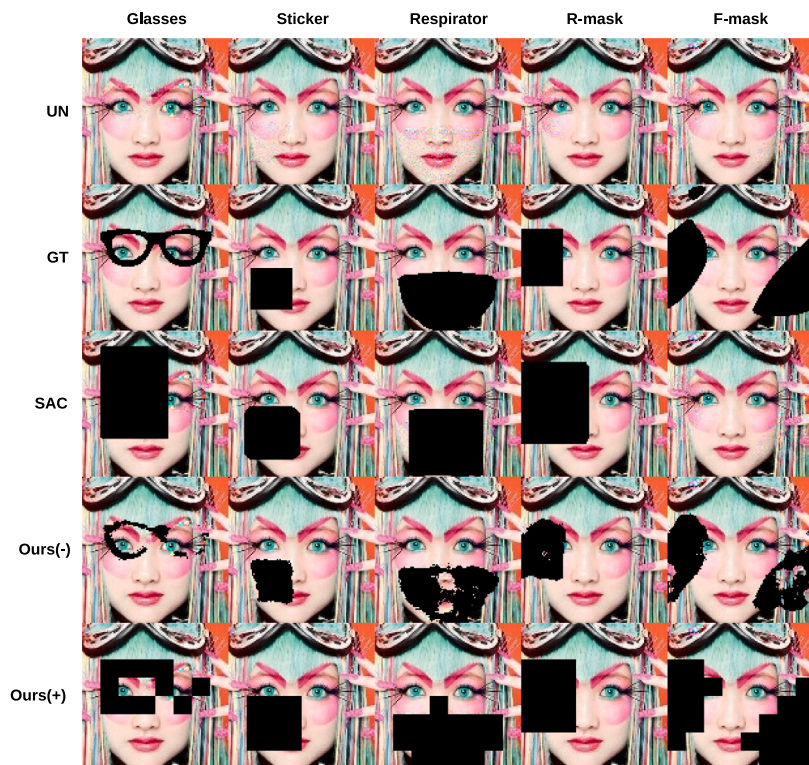
Target model	Defense method	Clean	Mask				
			Glasses	Sticker	Respirator	R-mask	F-mask
ArcMob	JPEG [15]	91.80	8.36	42.30	0.10	36.63	22.43
	LGS [10]	88.80	32.83	89.36	50.40	83.86	79.50
	SAC [18]	98.16	0.00	96.96	37.83	82.63	0.00
	<b>Ours(-)</b>	<b>98.20</b>	<b>83.93</b>	97.06	66.30	<b>92.40</b>	<b>90.16</b>
	<b>Ours(+)</b>	98.03	73.10	<b>97.10</b>	<b>66.60</b>	89.23	86.96
MagMob	JPEG [15]	89.63	8.70	43.80	0.23	34.03	25.36
	LGS [10]	91.33	29.43	90.53	57.03	83.50	79.90
	SAC [18]	<b>98.33</b>	0.66	97.50	56.06	85.93	0.00
	<b>Ours(-)</b>	98.30	<b>79.63</b>	<b>97.83</b>	<b>72.10</b>	<b>91.43</b>	<b>89.36</b>
	<b>Ours(+)</b>	98.10	76.03	97.60	66.90	89.30	87.20
MvMob	JPEG [15]	90.80	6.56	38.40	0.13	34.20	25.36
	LGS [10]	91.13	42.16	87.66	57.53	83.53	81.40
	SAC [18]	<b>98.33</b>	0.33	<b>97.23</b>	56.33	86.30	0.00
	<b>Ours(-)</b>	98.30	<b>85.13</b>	97.06	<b>70.10</b>	<b>92.80</b>	<b>89.86</b>
	<b>Ours(+)</b>	98.03	84.13	97.16	68.30	91.70	88.20
MvRes50	JPEG [15]	97.90	49.06	74.10	1.96	52.70	44.26
	LGS [10]	97.63	73.66	96.46	74.26	91.96	91.96
	SAC [18]	<b>99.53</b>	6.46	99.23	81.80	94.60	0.20
	<b>Ours(-)</b>	<b>99.53</b>	<b>97.93</b>	<b>99.30</b>	<b>92.80</b>	<b>97.53</b>	<b>97.50</b>
	<b>Ours(+)</b>	<b>99.53</b>	96.46	99.23	90.60	96.50	96.10
MvSwinT	JPEG [15]	98.20	91.00	94.46	76.73	90.86	91.96
	LGS [10]	99.30	97.10	99.46	92.86	96.96	97.03
	SAC [18]	99.60	2.23	<b>99.60</b>	85.76	94.60	0.00
	<b>Ours(-)</b>	<b>99.66</b>	<b>98.83</b>	99.50	<b>95.56</b>	<b>97.60</b>	<b>97.96</b>
	<b>Ours(+)</b>	99.53	97.73	99.46	93.20	96.70	96.36

**Table 10**  
Results of defense against adaptive attacks on closed-set FR systems. The attack method is PGDAP, the target model is MobileFaceNet, and the attacker’s goal is evasion.

Defense method	Clean	Mask				
		Glasses	Sticker	Respirator	R-mask	F-mask
DOA [9]	80.07	3.15	24.58	1.20	30.77	12.47
JPEG [15]	88.57	1.33	2.94	0.00	21.48	7.29
LGS [10]	87.59	18.09	25.11	0.00	28.47	9.53
PC [12]	92.70	5.18	63.21	0.44	49.30	12.10
SAC [18]	92.97	14.42	88.88	69.83	85.37	26.29
<b>Ours(-)</b>	94.00	16.45	71.90	4.61	61.38	26.40
<b>Ours(+)</b>	94.01	<b>82.85</b>	<b>91.15</b>	<b>81.81</b>	<b>89.76</b>	<b>69.75</b>



**Fig. 10.** Comparative results of non-adaptive and adaptive attacks.



**Fig. 11.** Visualization Results for Defending Against Adaptive Attacks.

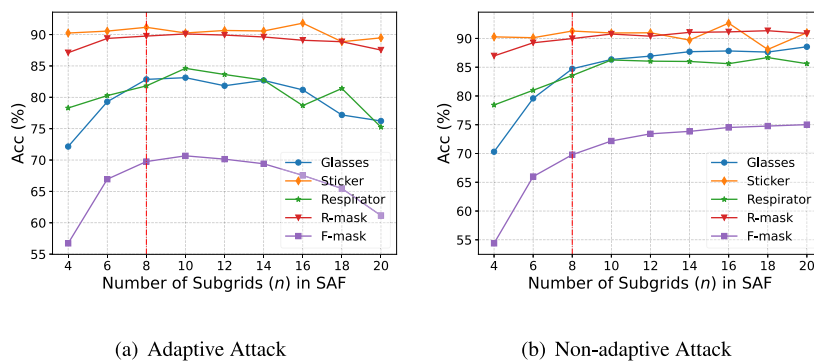


Fig. 12. Number of subgrids ( $n$ ) in SAF.

multiple holes. In contrast, “Ours(+)” effectively defends against both cases, showcasing the superior defense robustness and adaptability of our method.

In Fig. 12, we present defense results at different numbers of subgrids ( $n$ ). The figure indicates that the defense against adaptive attacks is influenced by the hyperparameter  $n$ . When  $n$  is too small, it tends to over-mask patches, while when  $n$  is too large, it under-masks patches. In our experiment, we set  $n$  to 8, but practical applications may require adjusting the hyperparameter  $n$  to optimize performance.

## 6. Discussion

**Different Patch Shapes.** PC [12] effectively defends against small patches but struggles with larger ones. Both PC and SAC [18] face challenges when dealing with intricate patch shapes. To enhance our method’s robustness and adaptability, we propose F-patch, which includes a variety of shapes for training the patch segmenter and the SAF strategy, enabling adaptive matching of patch shapes against patches of varying shapes.

**Multiple Patch Holes.** While both PC and SAC have mechanisms to address scenarios involving multiple patch holes, practical challenges arise when attempting to determine the exact number of patch holes employed by an attacker. As a result, their defense strategies prove less effective when dealing with situations involving random multiple patch holes. In contrast, our approach excels in providing adaptive defense in such complex cases.

**Obfuscated Gradients.** Similar to JPEG [15] and LGS [10], our method without SAF suffers from the problem of obfuscating the gradient. To mitigate this problem, we introduce the SAF strategy, which significantly enhances our defense method’s robustness.

**Limitation and Future Work.** Compared to theoretically provable defense methods like PC, incorporating a deep neural network structural patch segmenter in our approach presents challenges for theoretical proof. This limitation highlights an area for potential future work. Additionally, expanding the type and diversity of training data can further improve the robustness of adversarial patch recognition [37]. Furthermore, the RADAP can be enhanced to achieve automated detection of facial diseases from digital photographs with greater robustness and performance than recently proposed methods [38].

It is well known that images are generally noisy, and noise is a significant factor affecting the performance of automated recognition systems. Denoising methods are commonly applied to address this issue, and there are various techniques available. These methods should be employed carefully with appropriate parameters according to the characteristics of the images [39]. In future work, RADAP will be improved by adding a denoising module.

## 7. Conclusions

In conclusion, our work addresses the critical issue of defending against adversarial patches in both closed-set and open-set FR systems.

By introducing the RADAP method, we have presented a robust and adaptive defense approach capable of countering a wide range of adversarial patches. Our proposed techniques, FCutout and F-patch, utilizing randomly sampled masks in the Fourier space, significantly enhance the occlusion robustness of the FR model while optimizing the performance of the patch segmenter. Additionally, the introduction of the EBCE loss function further refines the patch segmenter’s capabilities. As an extension of this work, the effectiveness of the RADAP can be evaluated after employing it with a combined loss function designed with the EBCE or other combined loss functions such as used in [40]. To fortify our defense strategy against complete white-box adaptive attacks following the patch segmenter model’s leakage, we have introduced the SAF strategy.

The results of our extensive experiments not only validate the effectiveness of RADAP but also showcase its substantial improvements in defense performance against diverse adversarial patches. In fact, RADAP outperforms other state-of-the-art methods and even exhibits a higher clean accuracy compared to the undefended Vanilla model.

Our contributions underscore the significance of considering both closed-set and open-set systems in developing defense strategies for real-world FR applications. As adversarial threats continue to evolve, our work stands as a crucial step forward in enhancing the security and robustness of FR systems against adversarial attacks. We hope that our efforts will inspire further research in this vital domain and foster the development of even more resilient defense mechanisms.

## CRedit authorship contribution statement

**Xiaoliang Liu:** Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Furao Shen:** Supervision, Resources, Project administration, Investigation. **Jian Zhao:** Writing – review & editing. **Changhai Nie:** Supervision, Resources.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Data availability

The data that has been used is confidential.

## Acknowledgments

This work was supported in part by the STI 2030-Major Projects of China under Grant 2021ZD0201300, and by the National Science Foundation of China under Grant 62276127.

## References

- [1] I.J. Goodfellow, J. Shlens, C. Szegedy, Explaining and harnessing adversarial examples, 2014, arXiv preprint arXiv:1412.6572.
- [2] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, A. Vladu, Towards deep learning models resistant to adversarial attacks, in: International Conference on Learning Representations, ICLR, Learning Representations, 2018.
- [3] M. Sharif, S. Bhagavatula, L. Bauer, M.K. Reiter, Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition, in: ACM SIGSAC Conference on Computer and Communications Security, 2016, pp. 1528–1540.
- [4] X. Zheng, Y. Fan, B. Wu, Y. Zhang, J. Wang, S. Pan, Robust physical-world attacks on face recognition, *Pattern Recognit.* 133 (2023) 109009.
- [5] X. Liu, F. Shen, J. Zhao, C. Nie, RSTAM: An effective black-box impersonation attack on face recognition using a mobile and compact printer, 2022, arXiv preprint arXiv:2206.12590.
- [6] A. Zolfi, S. Avidan, Y. Elovici, A. Shabtai, Adversarial mask: Real-world universal adversarial attack on face recognition models, in: Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Springer, 2022, pp. 304–320.
- [7] P.-y. Chiang, R. Ni, A. Abdelkader, C. Zhu, C. Studer, T. Goldstein, Certified defenses for adversarial patches, in: International Conference on Learning Representations, ICLR, 2020.
- [8] Z. Zhang, B. Yuan, M. McCoyd, D. Wagner, Clipped bagnet: Defending against sticker attacks with clipped bag-of-features, in: 2020 IEEE Security and Privacy Workshops, SPW, IEEE, 2020, pp. 55–61.
- [9] T. Wu, L. Tong, Y. Vorobeychik, Defending against physically realizable attacks on image classification, in: International Conference on Learning Representations, ICLR, 2020.
- [10] M. Naseer, S. Khan, F. Porikli, Local gradients smoothing: Defense against localized adversarial attacks, in: 2019 IEEE Winter Conference on Applications of Computer Vision, WACV, IEEE, 2019, pp. 1300–1307.
- [11] C. Xiang, A.N. Bhagoji, V. Schwag, P. Mittal, {PatchGuard}: A provably robust defense against adversarial patches via small receptive fields and masking, in: 30th USENIX Security Symposium, USENIX Security 21, 2021, pp. 2237–2254.
- [12] C. Xiang, S. Mahloujifar, P. Mittal, {PatchCleanser}: Certifiably robust defense against adversarial patches for any image classifier, in: 31st USENIX Security Symposium, USENIX Security 22, 2022, pp. 2065–2082.
- [13] Z. Xiao, X. Gao, C. Fu, Y. Dong, W. Gao, X. Zhang, J. Zhou, J. Zhu, Improving transferability of adversarial patches on face recognition with generative models, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR, 2021, pp. 11845–11854.
- [14] X. Liu, F. Shen, J. Zhao, C. Nie, EAP: An effective black-box impersonation adversarial patch attack method on face recognition in the physical world, *Neurocomputing* 580 (2024) 127517.
- [15] G.K. Dziugaite, Z. Ghahramani, D.M. Roy, A study of the effect of jpg compression on adversarial images, 2016, arXiv preprint arXiv:1608.00853.
- [16] J. Hayes, On visible adversarial perturbations & digital watermarking, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, 2018, pp. 1597–1604.
- [17] Z. Qian, K. Huang, Q.-F. Wang, X.-Y. Zhang, A survey of robust adversarial training in pattern recognition: Fundamental, theory, and methodologies, *Pattern Recognit.* 131 (2022) 108889.
- [18] J. Liu, A. Levine, C.P. Lau, R. Chellappa, S. Feizi, Segment and complete: Defending object detectors against adversarial patch attacks with robust patch detection, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR, 2022, pp. 14973–14982.
- [19] E. Harris, A. Marcu, M. Painter, M. Niranjana, A. Prügell-Bennett, J. Hare, Fmix: Enhancing mixed sample data augmentation, 2020, arXiv preprint arXiv:2002.12047.
- [20] E. Goceri, Medical image data augmentation: techniques, comparisons and interpretations, *Artif. Intell. Rev.* 56 (2023) 12505–12561.
- [21] T. DeVries, G.W. Taylor, Improved regularization of convolutional neural networks with cutout, 2017, arXiv preprint arXiv:1708.04552.
- [22] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18, Springer, 2015, pp. 234–241.
- [23] A. Athalye, N. Carlini, D. Wagner, Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples, in: International Conference on Machine Learning, PMLR, 2018, pp. 274–283.
- [24] Q. Cao, L. Shen, W. Xie, O.M. Parkhi, A. Zisserman, Vggface2: A dataset for recognising faces across pose and age, in: 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition, FG 2018, IEEE, 2018, pp. 67–74.
- [25] Y. Guo, L. Zhang, Y. Hu, X. He, J. Gao, Ms-celeb-1m: A dataset and benchmark for large-scale face recognition, in: Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, the Netherlands, October 11–14, 2016, Proceedings, Part III 14, Springer, 2016, pp. 87–102.
- [26] G.B. Huang, M. Mattar, T. Berg, E. Learned-Miller, Labeled faces in the wild: A database for studying face recognition in unconstrained environments, in: Workshop on Faces in Real-Life Images: Detection, Alignment, and Recognition, 2008.
- [27] S. Chen, Y. Liu, X. Gao, Z. Han, Mobilefacenet: Efficient cnns for accurate real-time face verification on mobile devices, in: Biometric Recognition: 13th Chinese Conference, CCBR 2018, Urumqi, China, August 11–12, 2018, Proceedings 13, Springer, 2018, pp. 428–438.
- [28] J. Deng, J. Guo, N. Xue, S. Zafeiriou, Arcface: Additive angular margin loss for deep face recognition, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 4690–4699.
- [29] Q. Meng, S. Zhao, Z. Huang, F. 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, pp. 14225–14234.
- [30] X. Wang, S. Zhang, S. Wang, T. Fu, H. Shi, T. Mei, Mis-classified vector guided softmax loss for face recognition, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 34, 2020, pp. 12241–12248.
- [31] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
- [32] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, B. Guo, Swin transformer: Hierarchical vision transformer using shifted windows, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 10012–10022.
- [33] M. Tan, Q. Le, Efficientnet: Rethinking model scaling for convolutional neural networks, in: International Conference on Machine Learning, PMLR, 2019, pp. 6105–6114.
- [34] F. Croce, M. Hein, Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks, in: International Conference on Machine Learning, ICML, PMLR, 2020, pp. 2206–2216.
- [35] L. Bottou, Large-scale machine learning with stochastic gradient descent, in: Proceedings of COMPSTAT'2010, 2010, pp. 177–186.
- [36] I. Loshchilov, F. Hutter, SGDR: Stochastic gradient descent with warm restarts, in: International Conference on Learning Representations, ICLR, 2017.
- [37] M. Pintor, D. Angioni, A. Sotgiu, L. Demetrio, A. Demontis, B. Biggio, F. Roli, ImageNet-patch: A dataset for benchmarking machine learning robustness against adversarial patches, *Pattern Recognit.* 134 (2023) 109064.
- [38] E. Goceri, Convolutional neural network based desktop applications to classify dermatological diseases, in: The 4th IEEE International Conference on Image Processing, Applications and Systems, IPAS, 2020, pp. 138–143.
- [39] E. Goceri, Evaluation of denoising techniques to remove speckle and Gaussian noise from dermoscopy images, *Comput. Biol. Med.* 152 (2023) 106474.
- [40] E. Goceri, Polyp segmentation using a hybrid vision transformer and a novel loss function, *J. Digit. Imaging* 36 (2023) 1020–1034.

**Xiaoliang Liu** obtained his B.S. degree in communication engineering from Fujian University of Technology, Fujian, China. He is currently pursuing a Ph.D. in the Department of Computer Science and Technology at Nanjing University, Jiangsu, China. His research primarily focuses on adversarial machine learning, artificial neural networks, and computer vision.

**Furao Shen** (Member, IEEE) completed his B.Sc. and M.Sc. degrees in mathematics at Nanjing University, Nanjing, China, in 1995 and 1998, respectively. He earned his Ph.D. degree from the Tokyo Institute of Technology, Tokyo, Japan, in 2006. Currently, he holds the position of Full Professor of the School of Artificial Intelligence at Nanjing University. His research interests revolve around neural computing and robotic intelligence.

**Jian Zhao** (Senior Member, IEEE) earned his B.S. degree from Nanjing University, Nanjing, China, his M.Sc. degree from the Hamburg University of Technology, Hamburg, Germany, and his Dr. Sc. degree in electrical engineering from the Swiss Federal Institute of Technology (ETH) Zurich, Switzerland. From 2010 to 2015, he served as a Research Scientist at the Institute for Infocomm Research, A\*STAR, Singapore. Currently, he is an Associate Professor at the School of Electronic Science and Engineering, Nanjing University. His research interests encompass deep neural networks, mathematical optimization, and wireless communication networks. Dr. Zhao's achievements include the Dengfeng Scholars Program of Nanjing University in 2015, IEEE Globecom 2008 Best Paper Award, and the 2009 Chinese Government Award for Outstanding Self-Financed Students Abroad.

**Changhai Nie** holds the position of a professor in software engineering at the Department of Computer Science and Technology at Nanjing University, where he is affiliated with the National Key Laboratory for Novel Software Technology. His primary research interests lie in the fields of software testing and search-based software engineering, with a particular focus on combinatorial testing, search-based software testing, and methods for comparing and combining software testing techniques, among others.