

# Adv-Makeup: A New Imperceptible and Transferable Attack on Face Recognition

Bangjie Yin<sup>1\*</sup>, Wenxuan Wang<sup>2\*</sup>, Taiping Yao<sup>1\*†</sup>, Junfeng Guo<sup>3</sup>, Zelun Kong<sup>3</sup>,  
Shouhong Ding<sup>1†</sup>, Jilin Li<sup>1</sup> and Cong Liu<sup>3</sup>

<sup>1</sup>Youtu Lab, Tencent, Shanghai, China,

<sup>2</sup>Fudan University, Shanghai, China,

<sup>3</sup>The University of Texas at Dallas, Dallas, Texas, USA

{jamesyin10, wxwang.iris, gjf199509236817}@gmail.com, {taipingyao, ericshding, jerolinli}@tencent.com, {Zelun.Kong, cong}@utdallas.edu

## Abstract

Deep neural networks, particularly face recognition models, have been shown to be vulnerable to both digital and physical adversarial examples. However, existing adversarial examples against face recognition systems either lack transferability to black-box models, or fail to be implemented in practice. In this paper, we propose a unified adversarial face generation method - Adv-Makeup, which can realize imperceptible and transferable attack under the black-box setting. Adv-Makeup develops a task-driven makeup generation method with the blending module to synthesize imperceptible eye shadow over the orbital region on faces. And to achieve transferability, Adv-Makeup implements a fine-grained meta-learning based adversarial attack strategy to learn more vulnerable or sensitive features from various models. Compared to existing techniques, sufficient visualization results demonstrate that Adv-Makeup is capable to generate much more imperceptible attacks under both digital and physical scenarios. Meanwhile, extensive quantitative experiments show that Adv-Makeup can significantly improve the attack success rate under black-box setting, even attacking commercial systems.

## 1 Introduction

Deep neural networks (DNNs) are shown to achieve the state-of-the-art, and even human-competitive performance in many application domains, such as face recognition. However, recent works have proved that DNNs are vulnerable to adversarial examples [Goodfellow *et al.*, 2014; Dong *et al.*, 2018; Wang *et al.*, 2021], which are generated by adding perturbations over clean samples. These adversarial examples [Deng *et al.*, 2019; Dong *et al.*, 2019b] are effective against the state-of-the-art face recognition (FR) systems, which are pervasively adopted in a wide range of applications, from airport security check to mobile phone unlock payment.

\*Equal contributions.

†Corresponding author.

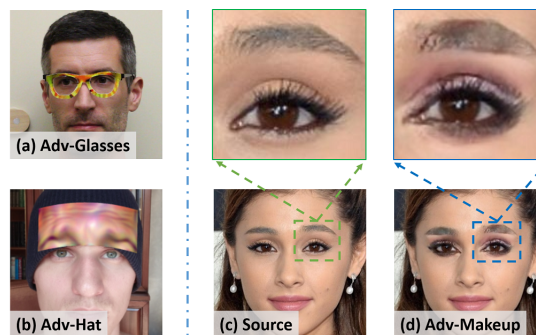


Figure 1: Adversarial faces generated by Adv-Glasses, Adv-Hat, and our Adv-Makeup. Impersonation attacks generated by Adv-Glasses and Adv-Hat are shown on the left. Source face, Adv-Makeup attack, and the corresponding zooming-in images are illustrated on the right.

Unfortunately, existing adversarial example generation methods on FR models exhibit several critical limitations that may severely constraint their applicability in practice. Specifically, we highlight the following limitations: (1) *Most existing methods are impractical and ineffective under physical scenarios.* A rich set of methods [Dong *et al.*, 2019b; Zhong and Deng, 2020; Song *et al.*, 2018; Deb *et al.*, 2019] could generate inconspicuous and strong perturbations over the entire facial images, focusing on the digital-world setting. Whereas, such global perturbations cannot be applied in realistic applications, because the attackers can only be able to manipulate their (physical) appearance, for example, it is impossible to add noise in the background area. (2) *Recent adversarial methods [Sharif *et al.*, 2016; Komkov and Petiushko, 2019; Brown *et al.*, 2017] which generate physical-implementable samples are rather noticeable which could be easily detected.* As illustrated in Fig.1 (a) and (b), the adversarial samples generated by such methods are visually obvious noticeable by human eyes, such as a pair of eyeglass, or imitating sticker located on the hat. (3) *The attacks are not transferable, i.e., not applicable to black-box FR models.* Existing methods mostly use the gradients of target model to construct adversarial examples. However, the target model is typically inaccessible directly in practical scenarios. Thus, it is notoriously hard for these attacks to successfully

fool the black-box FR model.

To tackle the above-discussed challenges simultaneously, we propose a novel imperceptible and transferable attack technique, Adv-Makeup, for face recognition. Firstly, the generated attacks shall be feasible in real-world applications. Adv-Makeup focuses on a common and practically implementable scenario: adding makeup to eye regions which shall mislead FR models yet being visually unnoticeable (*i.e.*, appearing as natural makeup). We develop a makeup generation module to synthesize realistic eye shadow on the source identity. To achieve imperceptible attack, we also apply a makeup blending method to alleviate the style and content differences between the source faces and generated eye shadows, which can further improve the stealthiness of the resulting adversarial faces. As shown in Fig.1 (c) and (d), the synthesized faces of Adv-Makeup with generated eye shadow appear natural and inconspicuous. Moreover, to enhance attack transferability, the adversarial examples should be model-agnostic. We thus introduce a fine-grained meta-learning attack strategy to learn better-generalized adversarial perturbations to improve the transferability of generated attack faces under the black-box scenario.

Our contributions are summarized as follows, (1) We propose a novel general Adv-Makeup technique, which can implement an imperceptible and transferable attack against face recognition models under both digital and physical scenarios. (2) To enhance the imperceptibility of attacks, we introduce a makeup generation module, which can add natural eye shadow over the orbital region. Moreover, a makeup blending strategy is proposed to ensure the consistency of style and content between source and generation, which can further improve the fidelity of the synthetic faces. (3) We also propose a task-driven fine-grained meta-learning adversarial attack strategy to guarantee the attacking effectiveness of the generated makeup, especially improving the transferability of adversarial examples on black-box victim models. (4) Through extensive and thorough experiments over the makeup dataset and LFW dataset, Adv-Makeup is proved to be able to generate imperceptible and transferable adversarial examples. A case study implementing Adv-Makeup in the physical world against two popular commercial FR platforms also demonstrates its efficacy in practice.

## 2 Related Work

**Adversarial Attacks.** Many attack algorithms have recently been proposed to show the deep learning models are broadly vulnerable to adversarial samples. For the white-box attack, there are lots of gradient-based approaches, such as [Goodfellow *et al.*, 2014; Carlini and Wagner, 2017; Madry *et al.*, 2017; Dong *et al.*, 2018]. However, such white-box attacking requires the adversarial examples to have full access to target models, which is unrealistic in real applications. Under a more practical scenario, there exist many black-box attack works [Wang *et al.*, 2021], where the adversary can not access the target model except only the outputs can be fetched. Whereas, the transferability of adversarial examples is the main challenge of the task. Even there already exist many studies trying to improve the transferabil-

ity [Dong *et al.*, 2019a; Xie *et al.*, 2019], the attack success rates of black-box attacks still decrease with a large margin. Moreover, the above attack methods apply perturbations to the entire input, which is impractical in the real world. In contrast, our Adv-Makeup method focuses on achieving implementable physical attack with high transferability.

**Adversarial Attacks on Face Recognition.** We mainly focus on the impersonation attack (targeted attack) on face recognition. As for the digital facial attacks, [Zhu *et al.*, 2019] firstly attempt to transfer eye-makeup to perform attacks on face recognition, and [Dong *et al.*, 2019b; Zhong and Deng, 2020; Sharif *et al.*, 2019; Song *et al.*, 2018; Deb *et al.*, 2019], however their adversarial examples either have obvious artifacts or can not directly apply in the physical world. Considering the wide application of FR in real-world applications, studies on physical attacks attract lots of researchers. [Sharif *et al.*, 2016] proposes a patch-based attack using optimization-based methods to add perturbations to the eyeglass region. Such attacks over the wearing hat are also generated in [Komkov and Petiushko, 2019]. However, these adversarial patches are perceptible to humans, and are mostly applied in the white-box setting, which is impractical in real-world scenarios. Different from prior works, our Adv-Makeup method can boost the similarity between the adversarial examples and targets under black-box scenario in both digital and physical conditions, and the synthesized makeup is imperceptible to be noticed.

## 3 Methodology

### 3.1 Overview

We propose a unified framework, *Adv-Makeup*, which can generate makeup-like perturbations over the orbital region of two eyes and learn better-generalized adversarial features from various models. As shown in Fig. 2, the Adv-Makeup framework comprises of three components, *i.e.*, makeup generation, makeup blending, and makeup attack. Taking the non-makeup source facial image and random face with cosmetics as inputs, the makeup generation can synthesize the face with realistic eye shadow. To further improve the visually-indistinguishable quality of generated adversarial samples, the makeup blending module helps to fuse the generated orbital region to the source face. Meanwhile, a fine-grained meta-learning adversarial attack module is introduced to make the generated faces adversarial as well as further improve the transferability of the attacks under black-box condition.

### 3.2 Makeup Generation

The goal of Adv-Makeup is to generate inconspicuous adversarial faces. Considering that eye-makeup is the most common facial alterations in daily life, owns various appearances, and can be implemented easily, besides, the eye area is a distinguishable region when performing face recognition, thus, we intend to firstly apply a makeup generation module to synthesize natural eye shadow over the orbital regions. To achieve this goal, we naturally turn to generative adversarial networks (GANs), which are widely used to synthesize images [Zhai *et al.*, 2019] and conduct makeup transfer [Wang

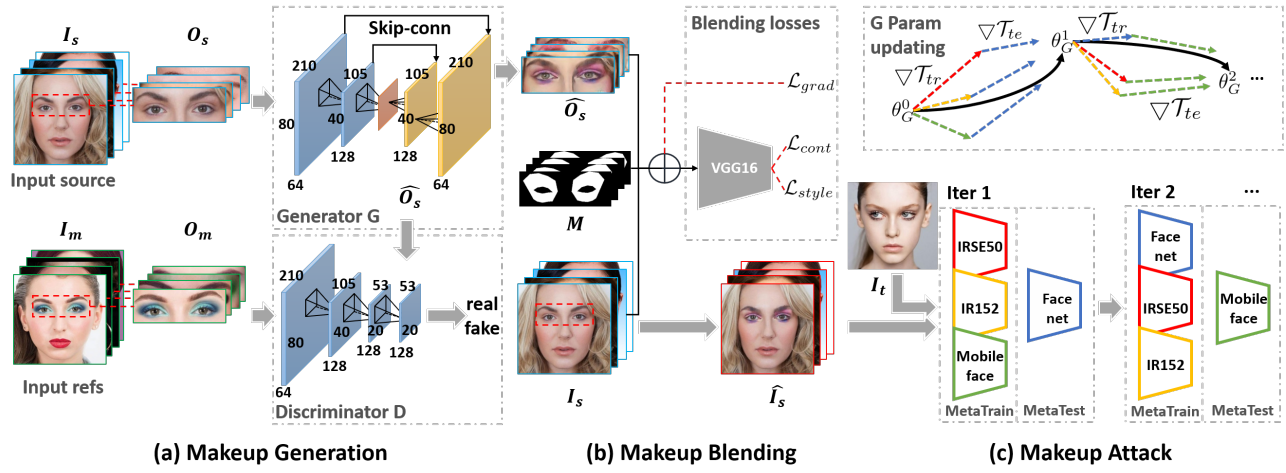


Figure 2: Overview of unified Adv-Makeup framework, which consists of makeup generation, makeup blending, and fine-grained meta-learning adversarial attack. (a) Makeup generation generates faces with eye shadow over the orbital regions. (b) Makeup blending is further applied to make the eye shadow more natural to achieve imperceptible generation. (c) Fine-grained meta-learning adversarial attack realize the transferable impersonation attack. Different colors in ‘G param updating’ denote various victim face recognition models.

*et al.*, 2020; Zhu *et al.*, 2019]. As shown in Fig. 2, the makeup generation contains a generator  $G$  to synthesize eye shadow, and a discriminator  $D$  to encourage perceptual realism of the generated images.

With the input of source image  $I_s \in D_s$  and makeup image  $I_m \in D_m$ , firstly we apply face alignment method [Zhang *et al.*, 2016] to generate corresponding landmarks for different faces. We align the faces according to affine matrix, which will also be used to align the landmarks. The landmarks surrounding two eyes are used for generating the 0-1 mask, and get the orbital regions  $O_s$  and  $O_m$ . Then the proposed generator takes orbital region  $O_s$  of the source image as input, and output  $\hat{O}_s = G(O_s)$  with generative eye shadow. Next, the resulting makeup orbital region can be calculated through  $O_s \odot (1 - M) + \hat{O}_s \odot M$ , where  $\odot$  is the element-wise product,  $M$  with equal-size of the  $O_s$  is a binary mask indicating the orbital region over the face. This region will be attached to the source face  $I_s$  to get the resulting face  $\hat{I}_s$ .

To improve the quality of the generation, we also introduce a discriminator encouraging the generated images to be more perceptually natural. The discriminator takes the real cosmetic orbit  $O_m$  of the makeup image and the output  $\hat{O}_s$  with synthesized eye shadow from generator as input, and it enforces the generation as natural as real eye shadow. The two-player ‘mini-max’ game makes the generation visually natural, so we define the generator loss  $\mathcal{L}_{gen}$  and discriminator loss  $\mathcal{L}_{dis}$  as,

$$\mathcal{L}_{gen} = E_{O_s}[\log(1 - D(G(O_s)))] \quad (1)$$

$$\mathcal{L}_{dis} = -[E_{O_m}[\log(D(O_m))] + E_{O_s}[\log(1 - D(G(O_s)))] \quad (2)$$

### 3.3 Makeup Blending

By applying such makeup generation, we can generate natural eye shadow over the orbital region, however, directly attaching the synthesized orbit region to the source image by

using mask would yield obvious style differences in content and abrupt intensity changes at the boundary. To eliminate the noticeable artifacts at the boundary and the style induced by the eye-shadow patches, we propose a makeup blending method to realize the imperceptible generation.

First, a gradient domain constraint is proposed, which we convert to a differentiable loss function, to alleviate the changes over generated boundary. By minimizing the loss function, the image details of generated face  $\hat{I}_s$  are able to be preserved while shifting the color to match the original image  $I_s$ . The gradient constraint loss  $\mathcal{L}_{grad}$  is defined as,

$$\mathcal{L}_{grad} = \|\nabla[I_s \odot (1 - M^*) + \nabla h(\hat{O}_s) \odot M^*] - \nabla \hat{I}_s\|_2^2 \quad (3)$$

where  $\nabla$  means the gradient operator, and  $M^*$  is a 0-1 mask expanded from  $M$  with the same size as  $I_s$ ,  $h$  attach  $\hat{O}_s$  to a all 0s matrix with the size of  $I_s$ .

Except for applying the  $\mathcal{L}_{grad}$  to control the smoothness and style over the  $\hat{I}_s$ , we also try to enhance the integration of style and content to better advance the naturalness of synthesized eye-shadow. Therefore, inspired by the [Gatys *et al.*, 2016], we utilize a pre-trained VGG16 model to calculate the style and content loss functions, *i.e.*,  $\mathcal{L}_{style}$  and  $\mathcal{L}_{cont}$ , pushing the style of eye-shadow patches closer to the source image and also preserving the content of the synthesized region. The content loss and style loss can be defined as,

$$\mathcal{L}_{cont} = \sum_{p=1}^P \frac{\alpha_p}{2N_p M_p} \sum_{j=1}^{N_p} \sum_{k=1}^{M_p} [(A_p[\hat{I}_s] - A_p[I_s]) \odot M^*]_{jk}^2 \quad (4)$$

$$\mathcal{L}_{style} = \sum_{p=1}^P \frac{\beta_p}{2N_p^2} \sum_{j=1}^{N_p} \sum_{k=1}^{N_p} (B_p[\hat{I}_s] - B_p[I_s])_{jk}^2 \quad (5)$$

where  $P$  is the number of convolutional VGG layers,  $N_p$  is the number of channels in activation,  $M_p$  is the number of

flattened activation values in each channel.  $A_p[\cdot] \in R^{N_p \times M_p}$  is an activation matrix computed from a deep network at corresponding layers.  $B_p[\cdot] = A_p[\cdot]A_p[\cdot]^T \in R^{N_p \times N_p}$  is the Gram matrix of features extracted at a set of VGG layers, and the  $\alpha_p$  and  $\beta_p$  are the weights to balance the contribution of each layer when calculating content and style loss.

### 3.4 Makeup Attack

**Adversarial Attack.** To make the generated makeup adversarial, we introduce an attack loss function against FR models. Leveraging a pre-trained face feature extractor  $F$ , the loss of impersonating attack ( $\mathcal{L}_{target}$ ) can be expressed as,

$$\mathcal{T} = 1 - \cos[F(I_t), F(\hat{I}_s)] \quad (6)$$

where  $I_t \in D_t$ , denoting the target image. The loss objective is utilized to make the source image wearing disturbed makeup  $\hat{I}_s$  to be close to the target image  $I_t$ . Nevertheless, for the black-box attack, such adversarial examples can not ‘generalize’ well on gradient-unknown models, resulting in low transferability. Recently, through finding the common-shared gradient direction among different victim models, ensemble training [Liu *et al.*, 2016; Dong *et al.*, 2018] can significantly improve the transferability of adversarial perturbations. In spite of that, ensemble training still has high chance to ‘overfit’ to the white-box models, thus perform ineffectively against the unknown victim models. Meta-learning has been applied in many research fields, like domain generation, and there are also some relevant applications in adversarial attacking/defence [Du *et al.*, 2019; Yin *et al.*, 2018]. However, they either focus on query-based black-box attack or enhancing network robustness. Inspired by [Shao *et al.*, 2020], we think meta-learning can also push the frontier of transferable attacks, therefore the fine-grained meta-learning adversarial attack is proposed to boost the black-box transferability.

**Meta-Train.** Assume we have  $L$  pre-trained victim FR models,  $\mathcal{F} = F_1, F_2, \dots, F_L$ , then we choose  $L - 1$  models to be meta-train models with performing impersonating attacks in each model,

$$\mathcal{T}_{tr}^i(\theta_G) = 1 - \cos[F_i(I_t), F_i(G(I_s))] \quad (7)$$

where  $i \in \{1, \dots, L - 1\}$  and  $\theta_G$  is the parameter of generator,  $G(I_s)$  indicates the faces with generated eye shadow from the makeup generator accompanying with makeup blending method.  $\mathcal{T}_{tr}^i(\theta_G)$  indicates that different models share a same generator, collecting the useful gradient information from each of the  $L - 1$  models to be a proper prior. To obtain  $i_{th}$  copy when updating  $\theta_G$ , we have  $\theta_G^i \leftarrow \theta_G - \alpha_1 \nabla_{\theta_G} \mathcal{T}_{tr}^i(\theta_G)$ .

**Meta-Test.** Moreover, we still have one remaining victim model,  $F_L$ , which will be the meta-test model. After obtaining the updated parameters, we conduct the impersonating attack by using the generator with  $\theta_G^i$  on  $F_L$ ,

$$\mathcal{T}_{te}^i(\theta_G^i) = 1 - \cos[F_L(I_t), F_L(G(I_s))] \quad (8)$$

---

#### Algorithm 1: The proposed Adv-Makeup.

---

**Require:**

**Input:** Source image  $I_s \in D_s$ ; target image  $I_t \in D_t$ ; makeup image  $I_m \in D_m$ ; pre-trained face models  $\mathcal{F} = F_1, F_2, \dots, F_L$ ; iterations  $T$ ; generator  $G$  and discriminator  $D$ .

**Initialization:** Model parameters  $\theta_G, \theta_D$ ; hyperparameters  $\alpha_1, \alpha_2, \beta_1, \beta_2, \beta_3$ .

**Ensure:** Model parameters  $\theta_G^*, \theta_D^*$ .

1: **for** each  $i \in T$  **do**

2: Update D: Compute  $\mathcal{L}_{dis}$ , then update

$$\theta_D' \leftarrow \theta_D - \nabla_{\theta_D} \mathcal{L}_{dis}(\theta_D)$$

3: Update G:

4: **Meta-train:**

5: Randomly select  $L - 1$  models from  $\mathcal{F}$  as meta-train models

6: Calculate  $(\mathcal{T}_{tr}^i, \mathcal{L}_{gen}, \mathcal{L}_{grad}, \mathcal{L}_{cont}, \mathcal{L}_{style})$  and

$$\text{then update } \theta_G^i \leftarrow \theta_G - \alpha_1 \nabla_{\theta_G} \mathcal{T}_{tr}^i(\theta_G)$$

7: **Meta-test:**

8: Use the remaining  $F_L$  to be meta-test model

9: Calculate  $\mathcal{T}_{te}^i$

10: **Meta-optimization:**

11: With Eq. 9

12: **endfor**

13:  $\theta_G^* = \theta_G'', \theta_D^* = \theta_D'$

14: **return**  $\theta_G^*, \theta_D^*$ ;

---

**Meta-Optimization.** To collect all the information in two stages, we propose a joint optimization strategy. The updating process can be written as Eq. 9, where  $\alpha_1$  and  $\alpha_2$  are the balancing weights for generation and attacking, and  $\beta_1, \beta_2$  and  $\beta_3$  are the corresponding parameters of makeup blending. Overall, the whole training process is illustrated in Alg.1.

$$\begin{aligned} \theta_G'' \leftarrow \theta_G - \alpha_1 \sum_{i=1}^{L-1} \nabla_{\theta_G} (\mathcal{T}_{tr}^i(\theta_G) + \mathcal{T}_{te}^i(\theta_G^i)) - \\ \alpha_2 \nabla_{\theta_G} \mathcal{L}_{gen}(\theta_G) - \beta_1 \nabla_{\theta_G} \mathcal{L}_{grad}(\theta_G) - \\ \beta_2 \nabla_{\theta_G} \mathcal{L}_{cont}(\theta_G) - \beta_3 \nabla_{\theta_G} \mathcal{L}_{style}(\theta_G) \end{aligned} \quad (9)$$

## 4 Experiment

### 4.1 Experimental Setup

**Datasets.** Two public datasets are utilized in the following experiments: 1) LFW [Huang *et al.*, 2008] contains 13233/5749 web-collected images/subjects with 6000 comparisons of same/different identities, most of them are low-quality. We randomly select 1000 different identity pairs to evaluate the attack performance. 2) [Gu *et al.*, 2019] released a high-quality makeup face database, including 333 frontal before-makeup faces and 302 after-makeup faces. By doing makeup generation, we select 198 after-makeup images with eye-makeup. In our digital attack setting, random 100 source before-makeup faces and 10 targets form 1000 comparisons for impersonation attacks. For both of the datasets, we report the average results across 1000 identity pairs.



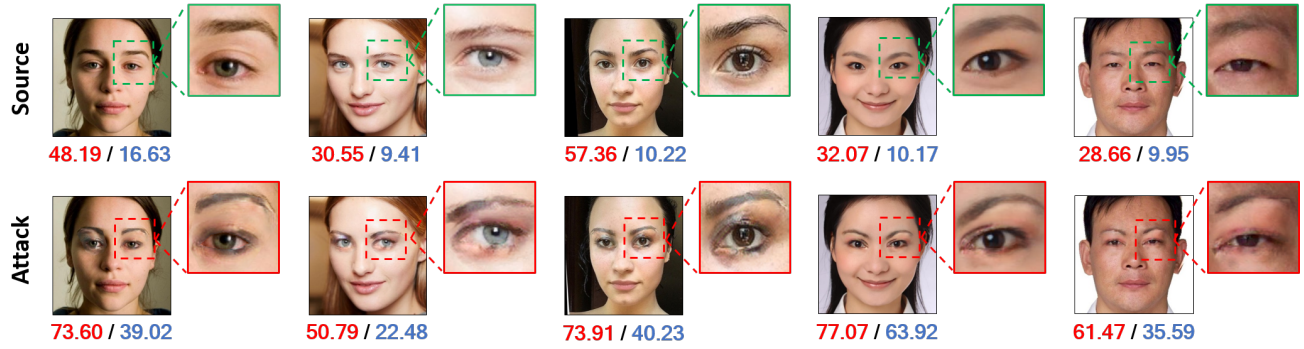


Figure 3: Generated adversarial faces from various illumination, pose, gender, and makeup intensity source faces. The target identity is the same target female in the first row of Fig.4. The *conf* results of each face resulting from commercial FR systems are pasted at the bottom. The red numbers before the slash are on Face++ and the blue number after the slash are from Microsoft.

**Competitors.** We choose several general attack methods, *i.e.*, FGSM [Goodfellow *et al.*, 2014], PGD [Madry *et al.*, 2017], MI-FGSM [Dong *et al.*, 2018], C&W [Carlini and Wagner, 2017]), and face attack works, *i.e.*, Adv-Glasses [Sharif *et al.*, 2016], Adv-Hat [Komkov and Petiushko, 2019], and Adv-Patch [Brown *et al.*, 2017], Adv-Makeup using ensemble training instead of fine-grained meta-learning (W.O. Meta in Tab.1), as the competing approaches. For a fair comparison, when comparing with the general attack methods, we restrict the attacking region only to the eye surroundings by multiplying the same 0-1 mask as ours.

**Implementation details.** The architectures of encoder and decoder are based on LADN [Gu *et al.*, 2019] with a U-Net structure. We employ an initial learning rate as 0.001 with momentum 0.9 via SGD. To balance different losses effects in Eq.9, we set  $\alpha_1, \alpha_2, \beta_1, \beta_2, \beta_3$  to be 1, 1, 0.1, 0.1, 0.1, respectively. We choose IR152, IRSE50, MobileFace [Deng *et al.*, 2019] and Facenet [Schroff *et al.*, 2015] as the victim FR models, 3 of them are selected for white-box training, and the remaining one is the black-box model. For all the comparison methods, we follow their official experimental settings.

**Attacking protocols.** The *attack success rate (ASR)* [Deb *et al.*, 2019; Zhong and Deng, 2020] is reported as the evaluation metric,

$$ASR = \frac{\sum_i^N 1_{\tau}(\cos[F(I_i^t), F(\hat{I}_s^t)] > \tau)}{N} \times 100\% \quad (10)$$

where  $1_{\tau}$  denotes the indicator function. We mainly consider impersonating attack, therefore the proportion of comparisons with similarity larger than  $\tau$  will be obtained as ASR. The value of  $\tau$  will be set as the threshold at 0.01 FAR (False Acceptance Rate) for each victim FR model as most face recognition works do, *i.e.*, IR152 (0.167), IRSE50 (0.241), MobileFace (0.302) and Facenet (0.409). Meanwhile, to further explore the reliability of attacks, we choose two cross-platform commercial state-of-the-art online face recognition systems as black-box victim systems to attack, *i.e.*, Face++ [MEGVII, 2021] and Microsoft [Microsoft, 2021]. Thus, we also report the *confidence score (conf)*, which means the confidence between attack image and target one, returned by the online systems in visualization results.

## 4.2 Digital-attack

To prove the effectiveness of the proposed Adv-Makeup on digital images, we design the multi-model attack experiments to present the performance improvement beyond the traditional ensemble adversarial training applied by other attacks.

**Comparison with other attack methods.** As mentioned before, we select seven different attack methods to compare, where training and testing paradigms strictly follow the original setting. As shown in Tab.1, each column represents the ASR results on the remaining black-box victim model for transferable testing. Comparing to other attacks, Adv-Makeup achieves the best transferable results on all victim models and significantly outperforms the competitors.

**Quality of generated adversarial faces.** To exhibit the imperceptibility of Adv-Makeup, we show six adversarial faces generated by Adv-Makeup in Fig.3, two numbers below each face image are the *conf*s given by the commercial online FR systems. The adversarial faces and corresponding *conf* results show that our method is robust to various interference factors such as illumination, pose, gender, and the intensity of eye makeup. We also compare the visual results with Adv-Hat, Adv-Glasses, and Adv-patch. In Fig.4, the examples show that the attacking makeup patches generated by Adv-Makeup own the most visually-indistinguishable appearance, while other forms need larger size of attacking area and look artificially and unnaturally to humans. Besides, we test the attacks on two commercial FR systems, and the comparison quantitative results again highlight the strength of Adv-Makeup in transferable black-box attack.

## 4.3 Physical-attack

Comparing to digital-attack, physical-attack seems more threatening for real-world applications. Thus, we utilize a simple tattoo paster to produce the Adv-Makeup attacks, and paste the tattoo eye shadow over attacker orbital region. Attacking two commercial FR platforms, we choose 20 identities, distributed evenly across gender and age, as personators, and randomly select 20 targets from makeup dataset.

**Adv-Makeup achieves the best physical attack results on the black-box commercial FR platforms.** In Fig.5, we

Methods	Dataset	LFW Dataset				Makeup Dataset			
	Target Model	IR152	IRSE50	FaceNet	MobileFace	IR152	IRSE50	FaceNet	MobileFace
Gradient-based	FGSM	2.51	5.54	1.75	5.27	7.32	2.13	9.44	32.1
	PGD	4.73	16.2	2.25	14.62	13.41	40.50	9.87	51.68
	MI-FGSM	4.75	16.21	2.12	14.30	13.43	40.92	9.73	52.36
	C&W	1.64	6.95	1.23	5.56	8.59	28.66	21.34	33.71
Patch-based	Adv-Hat	3.65	15.71	2.75	20.36	8.04	34.22	12.21	41.36
	Adv-Glasses	2.77	12.15	3.26	12.36	8.05	28.91	14.33	34.35
	Adv-Patch	1.23	9.35	4.65	10.32	5.93	24.09	14.46	31.11
Ours	W.O. Meta	5.23	15.91	5.06	19.55	21.43	56.17	32.61	61.62
	Adv-Makeup	<b>7.59</b>	<b>17.16</b>	<b>5.98</b>	<b>22.03</b>	<b>23.25</b>	<b>59.06</b>	<b>33.17</b>	<b>63.74</b>

Table 1: ASR results of digital impersonation attack over makeup and LFW datasets. There are four models used for the experiments, *i.e.*, IR152, IRSE50, MobileFace, and FaceNet. For each column, we regard the written model as the target black-box model, and the rest three models are used for training. The results are all from black-box attack setting.

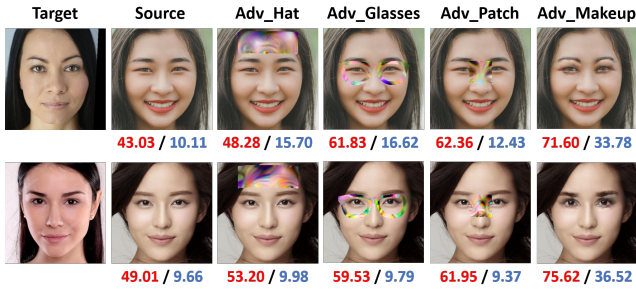


Figure 4: Comparison with other face attacks. The *conf* results of attacks to two black-box commercial FR platform are shown under the images. The red numbers before the slash are the *conf* results on Face++ and the blue numbers after the slash are from Microsoft.




Figure 6: Adversarial faces with **pose variations** under **physical-attack** against commercial FR platforms. The target identity is the same target female in the second row of Fig.4. The *conf* results of each attack are pasted at the bottom, the red numbers before the slash are on Face++, and the blue after the slash are from Microsoft.




Figure 5: The ASR comparison results of target attacks along with the confidence score threshold changes on Face++ and Microsoft.

show our average ASR results compared to Adv-Hat and Adv-Glasses according to different *conf* threshold returned by the platforms, and some visualization results are illustrated in Fig.6. It is clear that compared with current physical attacks, the Adv-Makeup attack strength is substantially higher than the others, which demonstrates our method can generate feasible transferable adversarial faces in real applications.

**Applying in real-world, the attacks of Adv-Makeup is the most imperceptible among the competitors.** As shown in Fig.6, we show harder attacker pairs with different target gender, to compare with other attacks. It is obvious that our changing area is the smallest, and adversarial faces are the most natural and inconspicuous.

**Adv-Makeup achieves better attack robustness to face pose changes.** In real-world scenes, most captured faces are not well-aligned, thus, the attack robustness to pose variations is crucial for physical attack. As in Fig. 6, ours achieves the most stable attack ability under large pose changes.

## 5 Conclusion

To implement a more powerful and feasible attack for real-world face recognition applications, we propose a novel method Adv-Makeup which can generate natural and seamless eye shadow to the orbital region, and achieve transferable attack under black-box models. The visualization results show that our method can generate much more imperceptible adversarial examples than other competitors. The experimental results on LFW and makeup datasets demonstrate the transferability of adversarial faces supervised by the task-driven fine-grained meta-learning adversarial attack strategy. The case study targeting two commercial FR platforms further verify the feasibility and efficacy of Adv-Makeup.

## References

- [Brown *et al.*, 2017] Tom B Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. *arXiv preprint arXiv:1712.09665*, 2017.
- [Carlini and Wagner, 2017] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *SP*, pages 39–57. IEEE, 2017.
- [Deb *et al.*, 2019] Debayan Deb, Jianbang Zhang, and Anil K Jain. Advfaces: Adversarial face synthesis. *arXiv preprint arXiv:1908.05008*, 2019.
- [Deng *et al.*, 2019] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019.
- [Dong *et al.*, 2018] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *CVPR*, pages 9185–9193, 2018.
- [Dong *et al.*, 2019a] Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *CVPR*, 2019.
- [Dong *et al.*, 2019b] Yinpeng Dong, Hang Su, Baoyuan Wu, Zhifeng Li, Wei Liu, Tong Zhang, and Jun Zhu. Efficient decision-based black-box adversarial attacks on face recognition. In *CVPR*, pages 7714–7722, 2019.
- [Du *et al.*, 2019] Jiawei Du, Hu Zhang, Joey Tianyi Zhou, Yi Yang, and Jiashi Feng. Query-efficient meta attack to deep neural networks. *arXiv preprint arXiv:1906.02398*, 2019.
- [Gatys *et al.*, 2016] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, pages 2414–2423, 2016.
- [Goodfellow *et al.*, 2014] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [Gu *et al.*, 2019] Qiao Gu, Guanzhi Wang, Mang Tik Chiu, Yu-Wing Tai, and Chi-Keung Tang. Ldn: Local adversarial disentangling network for facial makeup and demakeup. In *ICCV*, pages 10481–10490, 2019.
- [Huang *et al.*, 2008] Gary B. Huang, Marwan Mattar, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Technical Report*, 2008.
- [Komkov and Petiushko, 2019] Stepan Komkov and Aleksandr Petiushko. Advhat: Real-world adversarial attack on arcface face id system. *arXiv preprint arXiv:1908.08705*, 2019.
- [Liu *et al.*, 2016] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. *arXiv preprint arXiv:1611.02770*, 2016.
- [Madry *et al.*, 2017] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [MEGVII, 2021] MEGVII. Online face verification. <https://www.faceplusplus.com.cn/>, 2021.
- [Microsoft, 2021] Microsoft. Online face verification. <https://azure.microsoft.com/>, 2021.
- [Schroff *et al.*, 2015] Florian Schroff, Kalenichenko Dmitry, and Philbin James. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015.
- [Shao *et al.*, 2020] Rui Shao, Xiangyuan Lan, and Pong C. Yuen. Regularized fine-grained meta face anti-spoofing. In *AAAI*, pages 11974–11981, 2020.
- [Sharif *et al.*, 2016] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *CCS*, pages 1528–1540, 2016.
- [Sharif *et al.*, 2019] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K Reiter. A general framework for adversarial examples with objectives. *TOPS*, 2019.
- [Song *et al.*, 2018] Qing Song, Yingqi Wu, and Lu Yang. Attacks on state-of-the-art face recognition using attentional adversarial attack generative network. *arXiv preprint arXiv:1811.12026*, 2018.
- [Wang *et al.*, 2020] Wenxuan Wang, Yanwei Fu, Xuelin Qian, Yu-Gang Jiang, Qi Tian, and Xiangyang Xue. Fm2u-net: Face morphological multi-branch network for makeup-invariant face verification. In *CVPR*, 2020.
- [Wang *et al.*, 2021] Wenxuan Wang, Bangjie Yin, Taiping Yao, Li Zhang, Yanwei Fu, Shouhong Ding, Jilin Li, Feiyue Huang, and Xiangyang Xue. Delving into data: Effectively substitute training for black-box attack. *arXiv preprint arXiv:2104.12378*, 2021.
- [Xie *et al.*, 2019] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L Yuille. Improving transferability of adversarial examples with input diversity. In *CVPR*, pages 2730–2739, 2019.
- [Yin *et al.*, 2018] Chengxiang Yin, Jian Tang, Zhiyuan Xu, and Yanzhi Wang. Adversarial meta-learning. *arXiv preprint arXiv:1806.03316*, 2018.
- [Zhai *et al.*, 2019] Mengyao Zhai, Lei Chen, Frederick Tung, Jiawei He, Megha Nawhal, and Greg Mori. Lifelong gan: Continual learning for conditional image generation. In *ICCV*, pages 2759–2768, 2019.
- [Zhang *et al.*, 2016] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, 2016.
- [Zhong and Deng, 2020] Yaoyao Zhong and Weihong Deng. Towards transferable adversarial attack against deep face recognition. *arXiv preprint arXiv:2004.05790*, 2020.
- [Zhu *et al.*, 2019] Zheng-An Zhu, Yun-Zhong Lu, and Chen-Kuo Chiang. Generating adversarial examples by makeup attacks on face recognition. In *ICIP*. IEEE, 2019.