

Protecting Facial Privacy: Generating Adversarial Identity Masks via Style-robust Makeup Transfer

Shengshan Hu^{1, 4, 5, 6}, Xiaogeng Liu^{1, 4, 5, 6}, Yechao Zhang^{1, 4, 5, 6}, Minghui Li³

Leo Yu Zhang⁸, Hai Jin^{2, 4, 5, 7}, Libing Wu⁹

¹School of Cyber Science and Engineering, Huazhong University of Science and Technology

²School of Computer Science and Technology, Huazhong University of Science and Technology

³School of Software Engineering, Huazhong University of Science and Technology

⁴National Engineering Research Center for Big Data Technology and System

⁵Services Computing Technology and System Lab

⁶Hubei Engineering Research Center on Big Data Security

⁷Cluster and Grid Computing Lab ⁸School of Information Technology, Deakin University

⁹School of Cyber Science and Engineering, Wuhan University

{hushengshan, liuxiaogeng, ycz, minghuili, hjin}@hust.edu.cn

leo.zhang@deakin.edu.au, wu@whu.edu.cn

Abstract

While deep face recognition (FR) systems have shown amazing performance in identification and verification, they also arouse privacy concerns for their excessive surveillance on users, especially for public face images widely spread on social networks. Recently, some studies adopt adversarial examples to protect photos from being identified by unauthorized face recognition systems. However, existing methods of generating adversarial face images suffer from many limitations, such as awkward visual, white-box setting, weak transferability, making them difficult to be applied to protect face privacy in reality.

In this paper, we propose adversarial makeup transfer GAN (AMT-GAN)¹, a novel face protection method aiming at constructing adversarial face images that preserve stronger black-box transferability and better visual quality simultaneously. AMT-GAN leverages generative adversarial networks (GAN) to synthesize adversarial face images with makeup transferred from reference images. In particular, we introduce a new regularization module along with a joint training strategy to reconcile the conflicts between the adversarial noises and the cycle consistence loss in makeup transfer, achieving a desirable balance between the attack strength and visual changes. Extensive experiments verify that compared with state of the arts, AMT-GAN can not only preserve a comfortable visual quality, but also achieve a higher attack success rate over commercial FR APIs, including Face++, Aliyun, and Microsoft.

¹<https://github.com/CGCL-codes/AMT-GAN>




Figure 1. Comparison with existing adversarial attacks on FR systems in the black-box setting. The images in (a) are directly extracted from their papers. The numbers listed below images are the verification confidence of the target identity given by commercial FR APIs, and a higher score represent a stronger attack ability. The blue is from Face++, the green is from Aliyun, and the red is from Microsoft Azure (The same color code will be used hereinafter).

cluding Face++, Aliyun, and Microsoft.

1. Introduction

Recent years have witnessed the fast development of face recognition (FR) based on deep neural networks (DNNs).

The powerful face recognition systems, however, also pose a great threat to personal privacy. For example, it has been shown that FR systems can be used to identify social media profiles and track user relationships through large-scale photo analysis [16, 31]. Such kind of excessive surveillance on users urgently demands an effective approach to help individuals protect their face images against unauthorized FR systems.

Launching *data poisoning attacks* towards the training dataset or the gallery dataset of malicious FR models is a promising solution to protect facial privacy [3, 29]. However, these schemes require that the adversary (*i.e.*, the users in our scenario) can inject poisoned face images into the datasets. Once the target model is trained or makes inference over clean datasets, they are likely to become invalid.

Another strategy is to make use of *adversarial examples* to launch *evasion attacks* and protect face images from being illegally identified [26, 32, 40]. Adversarial face images are more suitable for protecting user privacy in real-world scenarios since they only need to modify users' own data regardless of the settings of the target model. Unfortunately, existing approaches of generating adversarial face examples suffer from several limitations when they are considered to protect face privacy on social media: (1) **Accessibility to target models**. Most existing schemes belong to white-box attack (*i.e.*, the adversary has full knowledge of the target model) [11, 24], or query-based black-box attack (*i.e.*, the adversary can arbitrarily query the target model) [9]. They are infeasible for protecting users' privacy since the users have no idea which kind of DNNs the third-party tracker is running; (2) **Poor visual quality**. As shown in Fig. 1, existing adversarial attacks on FR system fail to preserve the image quality in the black-box setting. Patch-based adversarial attacks [22, 37] often cause a fairly bizarre and conspicuous change on source images, and the state-of-the-art perturbation-based method [40] makes the modified face fill with awkward noises; (3) **Weak transferability**. Fig. 1 also demonstrates that the state of the art has a relatively low attack success rate on commercial APIs. *In summary, it is still challenging to balance the trade-off between the visual quality and attack ability of adversarial face images in the black-box setting.*

In this paper, we solve this problem from a new perspective. Different from existing works trying to place multifarious restrictions on perturbations and then dig a better gradient-based algorithm to construct adversarial examples, we focus on organizing the perturbations, although extensive and visible, in a reasonable way such that they appear natural and comfortable, under the condition that a high attack ability is maintained. We therefore leverage makeup as the key idea for arranging perturbations. Specifically, we propose a new framework called *adversarial makeup transfer GAN* (AMT-GAN) to generate adversarial face im-

ages with the natural appearance and stronger black-box attack strength. ATM-GAN first exploits a set of generative adversarial networks to construct adversarial examples that can inherit makeup styles from a reference image. In order to reconcile the conflicts between the adversarial noises and the cycle consistency loss in makeup transfer, we incorporate a newly designed regularization module by exploiting the disentanglement function of the encoder-decoder architecture and the residual dense blocks in the image super-resolution. As a result, the adversarial toxicity can be compatibly alleviated in the cycle reconstruction phase, making the generator focus on building robust mappings between the source domain and the target style domain with adversarial features. In addition, we introduce a joint training strategy which integrates the traditional *G-D* game in the GAN training and the newly designed regularization module, as well as the transferability enhancement process to encourage the generator to catch, imitate, and reconstruct the common adversarial features which can effectively transfer between different models.

To the best of our knowledge, we propose the first joint training framework to address the collapse phenomenon of the cycle consistency and the domain mappings of the generator when the image-to-image translation GANs are used to craft adversarial examples. Our joint training framework can be extended to other security-sensitive fields when GAN is considered, such as Deepfake [35]. In summary, we make the following contributions:

- We propose AMT-GAN, a more practical approach for protecting face images against unauthorized FR systems, by constructing adversarial examples with outstanding black-box attack performance and natural appearance that derive cosmetic styles from any chosen reference images.
- We design a regularization module based on feature disentanglement to improve the visual quality of adversarial images, and then develop a joint training pipeline to train the generator, the discriminator, and the regularization module, such that the generator can accomplish two jobs (*i.e.*, makeup-transfer and adversarial attack) simultaneously and build robust mappings among different data manifolds.
- Our extensive experiments on multiple benchmark datasets verify that AMT-GAN is highly effective at attacking various deep FR models, including commercial face verification APIs such as Face++², Aliyun³, and Microsoft⁴, where we outperform state of the arts about 4% ~ 60%.

²<https://www.faceplusplus.com>

³<https://vision.aliyun.com>

⁴<https://azure.microsoft.com>

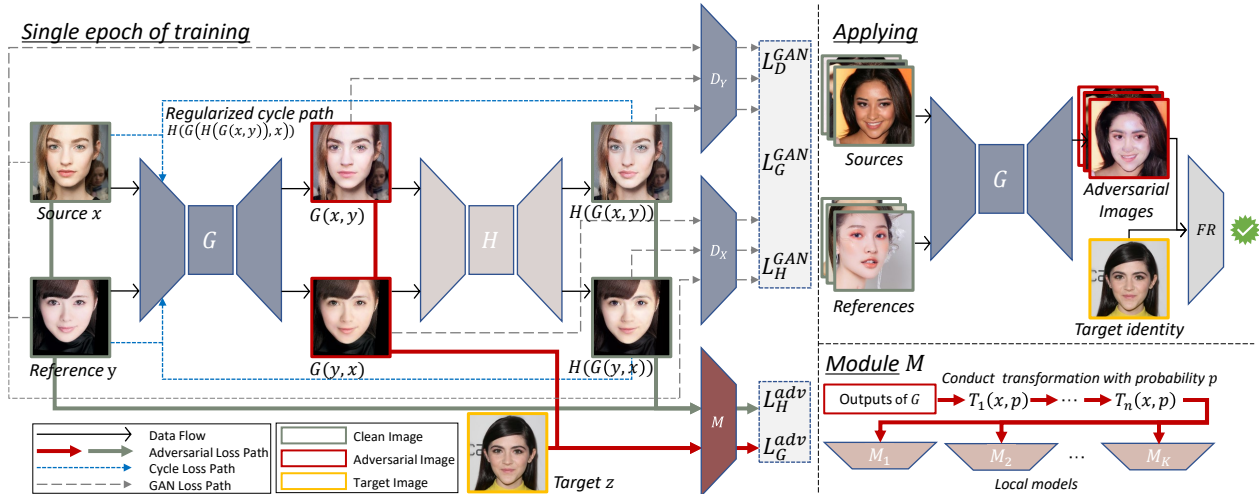


Figure 2. The architecture of AMT-GAN

2. Related Works

To protect face privacy, AMT-GAN leverages adversarial examples which are originally designed to disable DNNs in the machine learning community. Thus, we first discuss existing works on adversarial attacks in the context of privacy protection, followed by recent research on style transfer.

2.1. Adversarial Attacks on Face Recognition

Broad studies have shown that the DNNs are fatally vulnerable to adversarial examples [11, 24, 32], and many adversarial algorithms have been developed to attack face recognition system [9, 25, 30, 37]. Depending on the adversary’s knowledge of the target model, the adversarial face attacks can be divided into white-box attacks and black-box attacks.

The white-box attacks [11, 18, 24] and query-based black-box attacks [9, 13] heavily rely on the accessibility to the target model, which is a stringent prerequisite in practice. Transferability-based black-box attacks are thus more suitable for protecting face images in real-world scenarios [7, 37, 40, 44]. However, most existing transferability-based attacks [7, 38] are designed to solve an optimization problem, which is not only time-consuming, but also likely to be trapped in over-fitting and degrades the transferability [37]. In addition, to maintain the attack strength on different black-box models, transferability-based attacks usually generate images with perceptible noises [7, 8]. Recent work [40] tried to solve this problem by adding a new penalty function to fit the privacy-preserving scenario. However, as illustrated in Fig. 1, the perturbations are still outrageous. Using GAN [10] to craft adversarial samples [36] can achieve improvement in terms of image noises, but it is still challenging to maintain GAN’s stability for data with high-definition in adversarial task. [25]

turns to craft adversarial samples by swapping the latent features of face images from a GAN, however, it changes original face attributes dramatically (such as turning closed mouth to open), which are not desirable for users in social media. The patch-based adversarial attacks on FR systems [22, 30, 37, 41] are also not compatible for the same reason.

2.2. Style Transfer and Makeup Transfer

Style transfer [4] is an image-to-image translation technique that aims to separate and recombine the content and style information of images. Built on the style transfer framework, makeup transfer [1, 5, 12, 19, 23] is proposed to transfer the makeup style of the reference image to the source image while keeping the result of face recognition unchanged. Both style transfer and makeup transfer rely on the cycle consistency loss or its variants [43, 45] to maintain the stability of source images.

Recently, [46] made the first attempt to exploit the makeup transfer to generate adversarial face images in a white-box setting. Then [41] tried to construct adversarial cosmetic face images with transferability property to realize black-box attacks. However, it not only has a low attack success rate, but also fails to preserve the image visual quality where the modifications added to the source image is abnormal and noticeable, especially when the styles between reference and source images are significantly different, as explicitly depicted in our experiments Fig. 4.

It is indicated that adversarial noise may cause dysfunction on the cycle consistency loss [27, 39], and gets confirmed in this paper. In the literature, it is still challenging to amicably incorporate makeup transfer into adversarial examples, generating a natural adversarial face image that maintains a high attack success rate on black-box face

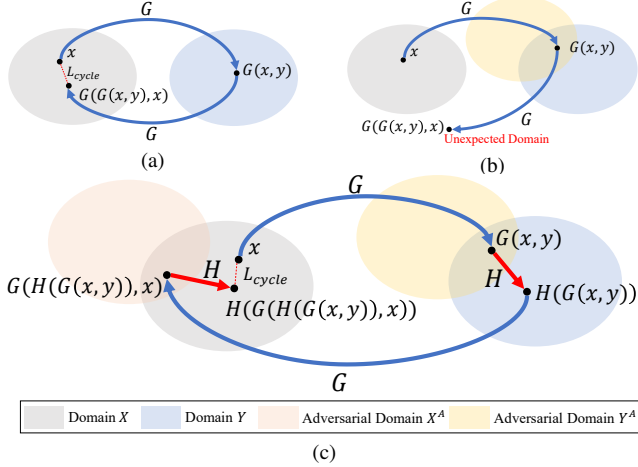


Figure 3. (a) Normal loop path of cycle consistency loss; (b) Damaged loop path caused by adversarial noises; (c) Recovered loop path of our regularized cycle consistency loss.

recognition systems.

3. Adversarial Makeup Transfer GAN (AMT-GAN)

3.1. Problem Formulation

In this section, we formulate the problem of adversarial attacks with makeup transfer on FR systems. To protect facial privacy effectively against malicious FR models, we mainly consider targeted adversarial attack (*i.e.* impersonation attack) that aims to generate adversarial examples which can be recognized as the specified target identity. Generally, the targeted adversarial attack on FR system can be formulated as:

$$\min_{x^A} L_{adv} = D(M_k(x^A), M_k(z)), \quad (1)$$

where $D(\cdot)$ represents a distance function such as cross-entropy or cosine similarity, M_k represents a DNN-based feature extractor for FR, x^A and z stand for the adversarial face image and the target image respectively.

As for makeup transfer, let $X, Y \subset \mathbb{R}^{H \times W \times 3}$ denote the makeup style domain of the source and reference images, respectively. Here, we use $x \in X$ and $y \in Y$ to represent the clean face images and $x^A \in X$ and $y^A \in Y$ to represent their adversarial face images, respectively. The adversarial makeup transfer is expected to train a function $G: \{x, y\} \rightarrow \tilde{y}_x^A$, where the adversarial image \tilde{y}_x^A has the same makeup style with y and the same visual identity with x .

3.2. Detailed Construction

The architecture of AMT-GAN is depicted in Fig 2.

Generator G and discriminators D_X, D_Y . The generator G is supposed to generate adversarial examples for

source images, while ensuring that the visual identity remains the same but the makeup style changes from the source domain to the reference domain. The discriminators D_X and D_Y are supposed to distinguish the distribution of fake images generated by G from that of the real images. Mathematically, the loss functions of GANs are formulated as:

$$L_D^{gan} = -\log D_X(x) - \log(1 - D_X(G(y, x))) - \log D_Y(y) - \log(1 - D_Y(G(x, y))), \quad (2)$$

$$L_G^{gan} = -\log(D_X(G(y, x))) - \log(D_Y(G(x, y))). \quad (3)$$

To learn the bidirectional mappings (*i.e.*, $X \rightarrow Y, Y \rightarrow X$) between the source and reference domain without supervised training data, we also utilize the cycle consistency loss function [45], which is a crucial element in the unsupervised image-to-image translation tasks. Generally, the cycle consistency loss function is depicted in Fig 3a and denoted as:

$$L^{cycle} = \|G(G(x, y), x) - x\|_1, \quad (4)$$

where $\|\cdot\|_1$ represents the L_1 norm.

However, the cycle consistency loss is in conflict with the adversarial example. As shown in Fig 3b, when establishing the inverse mapping $Y^A \rightarrow X$, the generated adversarial examples $G(x, y)$ is taken as the first input of G . Due to the adversarial modification on the samples from the clean domain X , G may fail to extract the features of $G(x, y)$ and is thus unable to turn an adversarial input from adversarial domain Y^A back to the clean domain X . As a result, the recovered $G(G(x, y), x)$ may be significantly different from the source image x , *i.e.*, $G(G(x, y), x)$ is in an unexpected domain rather than the domain X . In short, due to the adversarial nature of G 's outputs, it is hard for G to establish a robust inverse mapping $Y^A \rightarrow X$. This phenomenon makes the existing cycle consistency loss unsuitable for generating adversarial face images with makeup transfer (see our ablation study Fig. 6).

Regularization module H . To ensure the cycle consistency works well, we introduce a regularization module H in our framework. H is designed to generate clean images $H(G(x, y))$ with the same content, style, and dimensions as $G(x, y)$ but without adversarial property. Namely, H transforms the generated image $G(x, y)$ from the adversarial domain Y^A to the clean domain Y . The combination of G and H maintains a new cycle consistency, *i.e.*, the loop $X \rightarrow Y^A, Y^A \rightarrow Y$, then $Y \rightarrow X^A, X^A \rightarrow X$, as illustrated in Fig. 3c.

To find an effective H , we first choose a pair of encoder and decoder as the basic architecture of H . Then we leverage the *residual-in-residual dense block* (RRDB) as a key

block of H . RRDB is widely used in the field of image super-resolution [33]. These blocks will maintain and recover the content and texture information of the input while extracting and discarding the adversarial perturbations at the same time.

In our design, the regularization module H follows the same training process as the main networks G and D . The newly designed regularized cycle consistency loss is formulated as:

$$L_G^{reg} = \| H(G(H(G(x, y))), x) - x \|_1 + \| H(G(H(G(y, x))), y) - y \|_1. \quad (5)$$

In addition, H is supposed to keep the output of G visually unchanged while alleviating the adversarial effects. So we let H accept the penalty by the returns from D_X and D_Y , which encourages H to reconstruct the real performance of G in style transfer, and it is defined as:

$$L_H^{gan} = -\log(D_X(H(G(y, x)))) - \log(D_Y(H(G(x, y)))). \quad (6)$$

Note that in Eq. (2), we do not include the performance of H in L_D^{gan} , as we want to keep the two-player zero-sum game between the generator and the discriminators stable, and it is also unnecessary to adjust D_X and D_Y according to the performance of H .

Transferability enhancement module M . M consists of K pre-trained face recognition models $\{M_k\}_{k=1, \dots, K}$, which have high accuracy on public face images datasets. These local models serve as white-box models when we train our GANs and try to imitate the decision boundaries of potential target models which we cannot access.

In our method, inspired by [7, 38], we use an ensemble training strategy with input diversity enhancement to encourage G to generate adversarial examples with high transferability and black-box attack success rate. The adversarial is defined as:

$$L_G^{adv} = \frac{1}{2K} \sum_{k=1}^K 1 - \cos[M_k(z), M_k(T(G(x, y), p))] + \frac{1}{2K} \sum_{k=1}^K 1 - \cos[M_k(z), M_k(T(G(y, x), p))], \quad (7)$$

where M_k represents the feature extractor of the k -th local pre-trained white-box model, and we use cosine similarity as the distance function. $T(\cdot)$ represents the transformation function, and p is a pre-defined probability of whether the transformation will be conducted upon $G(x, y)$. Specifically, we choose image resizing and Gaussian noising as the transformation function. Both of them can degrade

Algorithm 1 The complete training process of AMT-GAN

Input: Source image set X ; reference image set Y ; target image z ; generator G ; regularization module H ; discriminators D_X, D_Y ; local models M ; optimizer $Adam$.

Parameter: Iterations T ; hyper-parameters Λ .

Output: parameters $\omega_G, \omega_{D_X}, \omega_{D_Y}, \omega_H$ for networks D_X, D_Y, G, H .

- 1: Initialize $\omega_G, \omega_{D_X}, \omega_{D_Y}, \omega_H$.
 - 2: **for** $i = 0$ to $T - 1$ **do**
 - 3: Randomly select source image $x \in X$ and reference image $y \in Y$ as the input of generator G ;
 - 4: **Updating D_X and D_Y with fixed G and H** ;
 - 5: Calculate L_D in Eq. (12);
 - 6: $\omega_{D_X} \leftarrow Adam(\omega_{D_X}, L_D)$;
 - 7: $\omega_{D_Y} \leftarrow Adam(\omega_{D_Y}, L_D)$;
 - 8: **Updating G with fixed D and H** ;
 - 9: Calculate L_G in Eq. (13);
 - 10: $\omega_G \leftarrow Adam(\omega_G, L_G)$;
 - 11: **Updating H with fixed G and D** ;
 - 12: Calculate L_H in Eq. (14);
 - 13: $\omega_H \leftarrow Adam(\omega_H, L_H)$;
 - 14: **end for**
 - 15: **return** $\omega_G, \omega_{D_X}, \omega_{D_Y}, \omega_H$.
-

the attack strength of adversarial examples whose adversarial modifications have faint transferability among different black-box models.

Accordingly, the adversarial attack loss of H is defined as:

$$L_H^{adv} = \frac{1}{2K} \sum_{k=1}^K 1 - \cos[M_k(x), M_k(H(G(x, y)))] + \frac{1}{2K} \sum_{k=1}^K 1 - \cos[M_k(y), M_k(H(G(y, x)))] \quad (8)$$

note that the input diversity is not included as it is unnecessary for H to own transferability.

Auxiliary Objectives. The histogram matching [23], denotes as $HM(x, y)$, is usually used to simulate the color distribution of reference y while preserves the content information of x . Here, we use this objective function to ensure the makeup similarity on lips, eye shadows, and face regions as well as the reconstruction ability of H . In detail, the makeup loss is defined as:

$$L_G^{make} = \| G(x, y) - HM(x, y) \|_2 + \| G(y, x) - HM(y, x) \|_2, \quad (9)$$

$$L_H^{make} = \| H(G(x, y)) - HM(x, y) \|_2 + \| H(G(y, x)) - HM(y, x) \|_2. \quad (10)$$

In addition, the generator G and the regularization module H are expected to preserve the original content and style information when the reference image is the source image itself, which is called self-reconstruction. This objective is significantly important for the generator G to maintain the structure information of resource images and avoid distortion of face attributes. The self-reconstruction path is defined as:

$$L_{G,H}^{idt} = \| H(G(x, x)) - x \|_1 + LPIPS(H(G(x, x)), x) \\ + \| H(G(y, y)) - y \|_1 + LPIPS(H(G(y, y)), y), \quad (11)$$

where the $LPIPS$ [42] function measures perceptual similarity between two images.

Total Loss. The total loss for D_X and D_Y is as follow:

$$L_D = L_D^{gan} \Lambda^T. \quad (12)$$

The total loss of G is defined as:

$$L_G = (L_G^{gan}, L_G^{reg}, L_G^{adv}, L_G^{make}, L_{G,H}^{idt}) \Lambda^T, \quad (13)$$

and the total loss of H is defined as:

$$L_H = (L_H^{gan}, L_H^{adv}, L_H^{make}, L_{G,H}^{idt}) \Lambda^T, \quad (14)$$

where $\Lambda = (\lambda_{gan}, \lambda_{reg}, \lambda_{adv}, \lambda_{make}, \lambda_{idt})$ represents the hyper-parameters. The entire training process is illustrated in Alg. 1.

4. Experiments

4.1. Experimental Setting

Implementation details. We construct the architecture of G , D_X , and D_Y in AMT-GAN following [19]. For the training process, the hyper-parameters λ_{GAN} , λ_{reg} , λ_{adv} , λ_{make} , and λ_{idt} are set to be 10, 10, 5, 2, and 5 respectively. We train the AMT-GAN by an Adam optimizer [21] with the learning rate of 0.0002, and set exponential decay rates as $(\beta_1, \beta_2) = (0.5, 0.999)$.

Competitors. We implement multiple benchmark schemes of adversarial attack, including PGD [24], MI-FGSM [7], TI-DIM [8], TIP-IM [40], and Adv-Makeup [41], to serve as the competitors for comparison. Note that PGD, MI-FGSM, and TI-DIM are very famous for their strong attack ability, TIP-IM is a very recent work which leverages adversarial examples to protect facial privacy, and Adv-makeup is the most relative scheme to ours which also exploits the makeup transfer to generate adversarial face images with transferability.

Datasets. Following [2, 19], the *Makeup Transfer* (MT) dataset [23] is used as the training dataset, which consists of 1115 non-makeup images and 2719 makeup images. We choose two datasets as our test sets: (1) CelebA-HQ [20] is

a widely used face image dataset with high quality. For the testing, we select a subset of CelebA-HQ, which contains 1,000 face images with different identities. (2) LADN-dataset [12] is a makeup dataset which contains 333 non-makeup images and 302 makeup images. We use 332 non-makeup images as the test images. We divide all the test images into 4 groups and aim images in each group to the same target identity.

Target models. Following [41], we conduct extensive experiments to attack 4 popular black-box FR models, which include IR152 [14], IRSE50 [17], Facenet [28], and Mobileface [6], and 3 commercial FR APIs including Face++, Aliyun, and Microsoft Azure.

Evaluation metrics. Following existing impersonation attacks [37, 41], we use *attack success rate* (ASR) to evaluate the attack ability of different methods. We calculate the ASR at FAR@0.01 for black-box testing. For commercial APIs, we directly record the confidence scores returned by FR servers. A higher confidence score represents that the victim FR API believes the two input images are of the same person with a higher probability. We also leverage FID [15], PSNR(dB), and SSIM [34] to evaluate the image quality. FID measures the distance between two data distributions, which is often used to investigate whether a generated dataset is as natural as the dataset extracted from the real world. PSNR and SSIM are widely-used methods to measure the difference between two images.

All of our experiments are conducted on RTX3090 GPU 24GB*1. For more implementation details and experimental results, please refer to our supplementary.

4.2. Comparison Study

Evaluations on black-box attacks. Tab. 1 shows black-box attacks on four different pre-trained models which have high accuracy on public datasets. For each target model, the other three models will serve as the ensemble training model. Note that for Adv-makeup these three models will serve as the meta-learning model. The results show that AMT-GAN has a strong attack ability in the black-box setting.

Evaluations on image quality. Tab. 2 shows the quantitative evaluations on image quality. Notably, compared with TIP-IM, although our method have a worse performance in terms of PSNR(db) and SSIM, AMT-GAN performs better for the FID result. This shows that the images generated by our method have more natural appearances than TIP-IM, although they get more information changed. This verifies our insight that arranging perturbations, instead of simply restricting them, is more important. We further attach the results of PSGAN [19], which is the most famous makeup transfer scheme, to show that it is normal to obtain similar evaluation results with these three metrics in the field of makeup transfer.

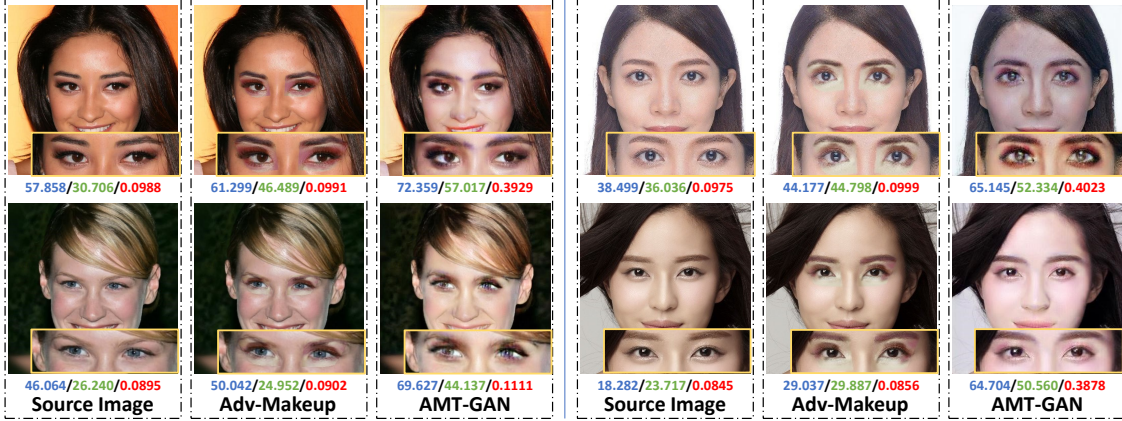


Figure 4. Comparison of visual quality between Adv-Makeup and AMT-GAN. The numbers under each image stand for the confidence scores returned from commercial APIs.

	CelebA-HQ				LADN-dataset			
	IRSE50	IR152	Facenet	Mobileface	IRSE50	IR152	Facenet	Mobileface
Clean	7.29	3.80	1.08	12.68	2.71	3.61	0.60	5.11
PGD [24]	36.87	20.68	1.85	43.99	40.09	19.59	3.82	41.09
MI-FGSM [7]	45.79	25.03	2.58	45.85	48.9	25.57	6.31	45.01
TI-DIM [8]	63.63	36.17	15.3	57.12	56.36	34.18	22.11	48.30
Adv-Makeup [41]	21.95	9.48	1.37	22.00	29.64	10.03	0.97	22.38
TIP-IM [40]	54.4	37.23	40.74	48.72	65.89	43.57	63.50	46.48
AMT-GAN	76.96	35.13	16.62	50.71	89.64	49.12	32.13	72.43

Table 1. Evaluations of *attack success rate* (ASR) for black-box attacks

In addition, Adv-makeup seems to behave well in all the quantitative evaluations. However, it has an extremely low attack success rate as demonstrated in Tab. 1. Furthermore, we give a qualitative comparison of visual image quality between Adv-Makeup and AMT-GAN, both of which construct adversarial face images based on makeup transfer. As shown in Fig. 4, the images generated by Adv-Makeup have sharp margins among the eyes region. On the contrary, AMT-GAN has a more realistic makeup style with smooth details. This is because Adv-makeup only changes the eyes region of original faces in a patch-based way, which leads to a good performance on quantitative evaluations, but leaves insufficient black-box attack strength and unresolved margin problem.

4.3. Attack Performance on Commercial APIs

Fig. 5 illustrates the attack performance towards Aliyun and Face++ for each test dataset. We collect and average the confidence scores from these APIs with massive adversarial examples. The results show that for both APIs, AMT-GAN outperforms competitors with regards to the attack ability.

4.4. Ablation Studies

Regularization module. Here we demonstrate the importance of the regularization module in maintaining the effectiveness and stability of AMT-GAN. As shown in Fig. 6

	FID(↓)	PSNR(↑)	SSIM(↑)
Adv-Makeup [41]	4.2282	34.5152	0.9850
TIP-IM [8]	38.7357	33.2089	0.9214
PSGAN [19]	27.6765	18.1403	0.8041
AMT-GAN (w/oH)	37.5486	19.3132	0.7807
AMT-GAN	34.4405	19.5045	0.7873

Table 2. Quantitative evaluations of image quality. AMT-GAN (w/oH) represents the AMT-GAN trained without the regularization module.

and Tab. 2, in the absence of the regularization module, the generator is likely to generate images with worse image quality, which indicates that the mappings between style domains are damaged to some degree. We owe this to the fact that the adversarial toxicity has poisoned the cycle reconstruction path. By applying the regularization module, the regularized cycle consistency loss can make the outputs of the generator more natural.

Style-robust makeup transfer. It is commonly expected that we can generate adversarial images with satisfied visual quality for any given makeup style. Thus it is desired to evaluate the impact of different references. We randomly choose 10 images from MT-dataset and LADN-dataset with different makeup styles as the references for testing. As illustrated in Fig. 7, AMT-GAN is robust to the changes of makeup style, where the targeted adversarial

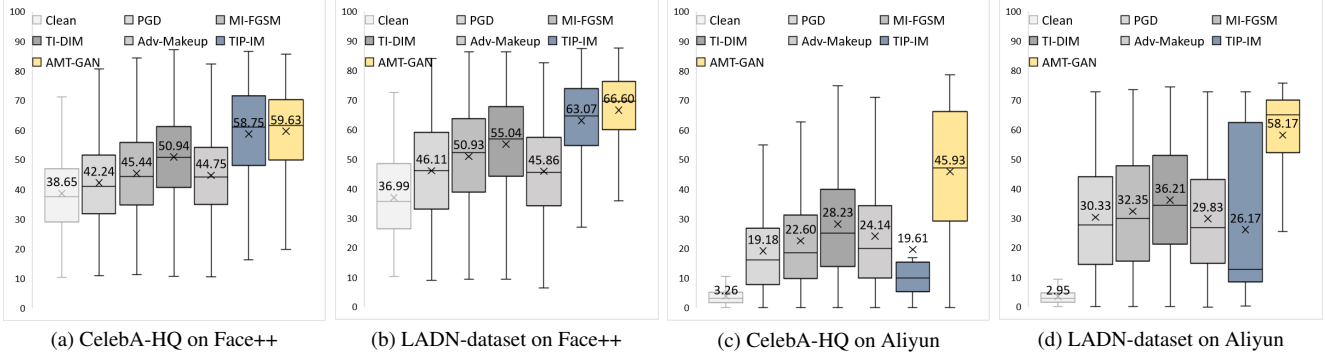


Figure 5. Confidence scores returned from Face++ and Aliyun. The mean confidence score of the runner-up (TIP-IM) for Face++ is 61.99, while ours is 64.58. For Aliyun, the mean confidence score of the runner-up (TI-DIM) is 32.37, while ours is 53.53. We outperform the runner-ups about 4% ~ 60%. Note that AMT-GAN also has a stronger transferability among different APIs while TIP-IM has a huge degradation in Aliyun compared with its performance on Face++.



Figure 6. Ablation study for the regularization module. The generator trained without the regularization module will generate images with unnatural details. Quantitative results are in Tab. 2.

face images maintain a good balance between the content of source images and the makeup style of references. The right figure of Fig. 7 shows that the changes of references have weak impact on the attack strength.




Figure 7. Evaluating the impact of different makeup styles. The left columns are adversarial examples under different references. The right figure illustrates average ASR towards different black-box models in Tab. 1 with 10 different references.

5. Limitations and Future Work

Although AMT-GAN shows effectiveness on attacking commercial APIs, it tends to have a higher attack strength and a better visual quality in images of female, which is caused by the unbalance of gender in makeup transfer training dataset (e.g., MT-dataset [23]). We believe that this problem can be solved by developing a more general and comprehensive training dataset.

Another problem with AMT-GAN is that the structure information sometimes gets slightly unaligned although we have designed corresponding objective functions for alleviation. The same problem also exists in the field of makeup

transfer [12, 19], and may become worse in generating adversarial examples. We leave this point to our future works.

Finally, the experimental results in Tab. 2 show that existing popular metrics used to evaluate the quality of images are unsuitable for the scenario of makeup transfer. New metrics are needed to evaluate whether a face image appears more natural than another one. Besides, it is still desirable for us to further improve the visual quality of adversarial faces. It is amazing to restrict the regions of makeup transfer to a small area (e.g., eyes), which can also preserve a high attack success rate. We also leave these to our future works.

6. Conclusion

In this paper, focusing on protecting facial privacy against malicious deep *face recognition* (FR) models, we propose AMT-GAN to construct adversarial examples that achieve a stronger attack ability in the black-box setting, while maintaining a better visual quality. AMT-GAN is able to generate adversarial face images with makeup transferred from any reference image. The experiments over multiple datasets and target models show that AMT-GAN is highly effective towards different open-source FR models and commercial APIs, and achieves a satisfied balance between the visual quality of adversarial face images and their attack strength.

Acknowledgments. Shengshan’s work is supported in part by the National Natural Science Foundation of China (Grant Nos. 62002126, U20A20177), and Fundamental Research Funds for the Central Universities (Grant No. 5003129001). Leo’s work is supported in part by the National Natural Science Foundation of China (Grant No. 61702221). Libing’s work is supported by Key R&D plan of Hubei Province (No. 2021BAA025). Minghui Li is the corresponding author.

References

- [1] Huiwen Chang, Jingwan Lu, Fisher Yu, and Adam Finkelstein. Pairedcyclegan: Asymmetric style transfer for applying and removing makeup. In *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'18)*, pages 40–48, 2018. 3
- [2] Hung-Jen Chen, Ka-Ming Hui, Szu-Yu Wang, Li-Wu Tsao, Hong-Han Shuai, and Wen-Huang Cheng. Beautyglow: On-demand makeup transfer framework with reversible generative network. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'19)*, pages 10042–10050, 2019. 6
- [3] Valeriia Cherepanova, Micah Goldblum, Harrison Foley, Shiyuan Duan, John P. Dickerson, Gavin Taylor, and Tom Goldstein. Lowkey: Leveraging adversarial attacks to protect social media users from facial recognition. In *Proceedings of the 9th International Conference on Learning Representations (ICLR'21)*, 2021. 2
- [4] Yunjey Choi, Min-Je Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'18)*, pages 8789–8797, 2018. 3
- [5] Han Deng, Chu Han, Hongmin Cai, Guoqiang Han, and Shengfeng He. Spatially-invariant style-codes controlled makeup transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6549–6557, June 2021. 3
- [6] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'19)*, pages 4690–4699, 2019. 6
- [7] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'18)*, pages 9185–9193, 2018. 3, 5, 6, 7
- [8] Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'19)*, pages 4312–4321, 2019. 3, 6, 7
- [9] 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 *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'19)*, pages 7714–7722, 2019. 2, 3
- [10] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial networks. *CoRR*, abs/1406.2661, 2014. 3
- [11] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR'15)*, 2015. 2, 3
- [12] Qiao Gu, Guanzhi Wang, Mang Tik Chiu, Yu-Wing Tai, and Chi-Keung Tang. LADN: local adversarial disentangling network for facial makeup and de-makeup. In *Proceedings of the 17th IEEE/CVF International Conference on Computer Vision (ICCV'19)*, pages 10480–10489, 2019. 3, 6, 8
- [13] Ying Guo, Xingxing Wei, Guoqiu Wang, and Bo Zhang. Meaningful adversarial stickers for face recognition in physical world. *CoRR*, abs/2104.06728, 2021. 3
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the 2016 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'16)*, pages 770–778, 2016. 6
- [15] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6626–6637, 2017. 6
- [16] Kashmir Hill. The secretive company that might end privacy as we know it. *The New York Times*, 18:2020, 2020. 2
- [17] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'18)*, pages 7132–7141, 2018. 6
- [18] Shengshan Hu, Yechao Zhang, Xiaogeng Liu, Leo Yu Zhang, Minghui Li, and Hai Jin. Advhash: Set-to-set targeted attack on deep hashing with one single adversarial patch. In *Proceedings of the 29th ACM International Conference on Multimedia (MM'21)*, pages 2335–2343, 2021. 3
- [19] Wentao Jiang, Si Liu, Chen Gao, Jie Cao, Ran He, Jiashi Feng, and Shuicheng Yan. PSGAN: pose and expression robust spatial-aware GAN for customizable makeup transfer. In *Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'20)*, pages 5193–5201, 2020. 3, 6, 7, 8
- [20] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 6
- [21] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR'15)*, 2015. 6
- [22] Stepan Komkov and Aleksandr Petiushko. Advhat: Real-world adversarial attack on arcface face ID system. In *Proceedings of the 25th International Conference on Pattern Recognition (ICPR'20)*, pages 819–826, 2020. 2, 3
- [23] Tingting Li, Ruihe Qian, Chao Dong, Si Liu, Qiong Yan, Wenwu Zhu, and Liang Lin. Beautygan: Instance-level facial makeup transfer with deep generative adversarial network. In *Proceedings of the 26th ACM International Conference on Multimedia (MM'18)*, pages 645–653, 2018. 3, 5, 6, 8
- [24] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *Proceedings of the*

- 6th International Conference on Learning Representations (ICLR'18), 2018. 2, 3, 6, 7
- [25] Haonan Qiu, Chaowei Xiao, Lei Yang, Xinchen Yan, Honglak Lee, and Bo Li. Semanticadv: Generating adversarial examples via attribute-conditioned image editing. In *Proceedings of the 16th European Conference on Computer Vision (ECCV'20)*, pages 19–37, 2020. 3
- [26] Arezoo Rajabi, Rakesh B. Bobba, Mike Rosulek, Charles V. Wright, and Wu-chi Feng. On the (im)practicality of adversarial perturbation for image privacy. *Proc. Priv. Enhancing Technol.*, 2021(1):85–106, 2021. 2
- [27] Nataniel Ruiz, Sarah Adel Bargal, and Stan Sclaroff. Disrupting deepfakes: Adversarial attacks against conditional image translation networks and facial manipulation systems. In *Proceedings of the 15th European Conference on Computer Vision Workshops (ECCVW'20)*, pages 236–251, 2020. 3
- [28] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the 2015 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'15)*, pages 815–823, 2015. 6
- [29] Shawn Shan, Emily Wenger, Jiayun Zhang, Huiying Li, Haitao Zheng, and Ben Y. Zhao. Fawkes: Protecting privacy against unauthorized deep learning models. In *Proceedings of the 29th USENIX Security Symposium (USENIX Security'20)*, pages 1589–1604, 2020. 2
- [30] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K. Reiter. Adversarial generative nets: Neural network attacks on state-of-the-art face recognition. *CoRR*, abs/1801.00349, 2018. 3
- [31] Yan Shoshitaishvili, Christopher Kruegel, and Giovanni Vigna. Portrait of a privacy invasion. *Proc. Priv. Enhancing Technol.*, 2015(1):41–60, 2015. 2
- [32] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR'14)*, 2014. 2, 3
- [33] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. ESRGAN: enhanced super-resolution generative adversarial networks. In *Proceedings of the 15th European Conference on Computer Vision Workshops (ECCVW'18)*, 2018. 5
- [34] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, 2004. 6
- [35] Mika Westerlund. The emergence of deepfake technology: A review. *Technology Innovation Management Review*, 9(11), 2019. 2
- [36] Chaowei Xiao, Bo Li, Jun-Yan Zhu, Warren He, Mingyan Liu, and Dawn Song. Generating adversarial examples with adversarial networks. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18)*, pages 3905–3911, 2018. 3
- [37] Zihao Xiao, Xianfeng Gao, Chilin Fu, Yinpeng Dong, Wei Gao, Xiaolu Zhang, Jun Zhou, and Jun Zhu. Improving transferability of adversarial patches on face recognition with generative models. In *Proceedings of the 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'21)*, pages 11845–11854, 2021. 2, 3, 6
- [38] 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 *Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'19)*, pages 2730–2739, 2019. 3, 5
- [39] Chaofei Yang, Leah Ding, Yiran Chen, and Hai Li. Defending against gan-based deepfake attacks via transformation-aware adversarial faces. In *Proceedings of the 2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2021. 3
- [40] Xiao Yang, Yinpeng Dong, Tianyu Pang, Hang Su, Jun Zhu, Yuefeng Chen, and Hui Xue. Towards face encryption by generating adversarial identity masks. In *Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV'21)*, pages 3897–3907, 2021. 2, 3, 6, 7
- [41] Bangjie Yin, Wenxuan Wang, Taiping Yao, Junfeng Guo, Zelin Kong, Shouhong Ding, Jilin Li, and Cong Liu. Advmakeup: A new imperceptible and transferable attack on face recognition. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI'21)*, pages 1252–1258, 2021. 3, 6, 7
- [42] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'18)*, pages 586–595, 2018. 6
- [43] Yihao Zhao, Ruihai Wu, and Hao Dong. Unpaired image-to-image translation using adversarial consistency loss. In *Proceedings of the 16th European Conference on Computer Vision (ECCV'20)*, volume 12354 of *Lecture Notes in Computer Science*, pages 800–815, 2020. 3
- [44] Yaoyao Zhong and Weihong Deng. Towards transferable adversarial attack against deep face recognition. *IEEE Trans. Inf. Forensics Secur.*, 16:1452–1466, 2021. 3
- [45] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV'21)*, pages 2242–2251, 2017. 3, 4
- [46] Zheng-An Zhu, Yun-Zhong Lu, and Chen-Kuo Chiang. Generating adversarial examples by makeup attacks on face recognition. In *Proceedings of the 2019 IEEE International Conference on Image Processing (ICIP'19)*, pages 2516–2520, 2019. 3