

On Adversarial Patches: Real-World Attack on ArcFace-100 Face Recognition System

Mikhail Pautov
Intelligent Systems Lab
Huawei Moscow Research Center
 Moscow, Russia
 mikhail.pautov@phystech.edu

Grigorii Melnikov
Intelligent Systems Lab
Huawei Moscow Research Center
 Moscow, Russia
 grigorii.melnikov@skoltech.ru

Edgar Kaziakhmedov *Intelligent Systems Lab*
Huawei Moscow Research Center
 Moscow, Russia
 edgar.kaziakhmedov@phystech.edu

Klim Kireev
Intelligent Systems Lab
Huawei Moscow Research Center
 Moscow, Russia
 klim.kireev@skoltech.ru

Aleksandr Petiushko
Intelligent Systems Lab
Huawei Moscow Research Center
 Moscow, Russia
 petyushko.alexander1@huawei.com

Abstract—Recent works showed the vulnerability of image classifiers to adversarial attacks in the digital domain. However, the majority of attacks involve adding small perturbation to an image to fool the classifier. Unfortunately, such procedures can not be used to conduct a real-world attack, where adding an adversarial attribute to the photo is a more practical approach. In this paper, we study the problem of real-world attacks on face recognition systems. We examine security of one of the best public face recognition systems, LResNet100E-IR with ArcFace loss, and propose a simple method to attack it in the physical world. The method suggests creating an adversarial patch that can be printed, added as a face attribute and photographed; the photo of a person with such attribute is then passed to the classifier such that the classifier's recognized class changes from correct to the desired one. Proposed generating procedure allows projecting adversarial patches not only on different areas of the face, such as nose or forehead but also on some wearable accessory, such as eyeglasses.

Index Terms—adversarial patch, face recognition, physical domain, ArcFace.

I. INTRODUCTION

In the last years, much of research was done in the field of attacks on face recognition systems. Nowadays, one of the best face recognition systems to test different attacking approaches is LResNet100E-IR with special additive angular margin loss, ArcFace. The network maps an image to a feature vector such that intra-class distance tends to be small and inter-class distance stays large. In the original paper [1] it is observed that the accuracy in classification problems of this network is comparable to state-of-the-art models.

Although deep neural networks are efficient in image classification, they are vulnerable to adversarial samples [2]. Majority of prior works has focused on exploiting such vulnerabilities through imperceptible changes of an image, which are constrained by some norm, e.g. L_1 , L_2 , L_∞ . In such cases, an attacker is trying to fool classifier but at the same time keeps perceptual similarity between original and adversarial images. Later the authors of [3] succeeded in localizing of

image changes. They generated the patch that can be placed anywhere within the field of view of a classifier and managed to conduct an attack in the real world.

Attack on face recognition neural network is more complex in comparison to attack on just image classifier. It is known for a long time that several areas on a face influence classification more than the others [4]. So, the result of classification of the photo with an adversarial patch depends on an area where the patch is located.

Authors of [5] investigated the possibility to construct an attack inconspicuous for the human eye. They managed to conduct an attack in the real world with the use of eyeglasses as an adversarial attribute. These glasses were constructed to impersonate a person, but were designed for the face of attacker with a fixed orientation of his head.

The goal of this paper is to make a starting point in the improvement of the security of the open face recognition system, ArcFace. Our approach is to explore its vulnerability to adversarial attacks and mislead the network using the proposed method. It should be mentioned that the network and its weights are available on the Internet and the adversarial attack approach is considered as well-known, so the research conducted does not violate any regulations.

The outline of current research is presented below:

- A simple adversarial patch generation procedure is proposed;
- The procedure allows to project adversarial patch on different surfaces on a face;
- Effectiveness of adversarial patches placed on different positions on the face was studied;
- A real-world attack on one of the best public face recognition systems with the use of this procedure is conducted;
- The attacking pipeline can be easily implemented since the gained patch is a gray-scale image.

II. THE MAIN CONCEPT AND RELATED WORKS

In this section, we describe a concept of adversarial attacks and present an overview of methods and techniques which are often used in the physical domain.

A. The concept of adversarial attacks

An adversarial attack on face recognition is a technique to fool some recognition system through a change in input such that the output of the system changes from correct to another one. Although adversarial attack may be conducted in the digital domain (where the input to the classifier may be changed, for example, pixel-wise), it is much harder to construct such an attack in the physical world. However, in [5], [6], it was shown that it is possible with the efficiency comparable to the one in the digital domain.

There are several possible classifications of adversarial attacks on face recognition systems (i.e. on classifiers). One way is to determine how the output class should change:

- targeted attack – the adversary changes the output classification of input to the desired one;
- untargeted attack – the adversary leads to misclassification of input.

In [7], the classification of attacks based on what is known about the neural network is provided. In white-box attacks, parameters of the model, as well as its structure and training procedure, are known. In contrast, in black-box scenarios, none of the above is known.

B. Related works

In recent years, much research had been done in the field of attacks in the physical domain.

In [8], one of the most straightforward attacking approaches, Fast Gradient Sign Method, FGSM, was proposed. This method refers to generating of adversarial examples through adding to the initial image the perturbation

$$\eta = \varepsilon \text{sign}(\nabla_x J(\theta, x, y)), \quad (1)$$

where θ is vector of parameters of the model, x is input to the model, y are the targets associated with x and $J(\theta, x, y)$ is the cost function used to train the neural network. Note that in case of targeted attack, y are desired targets and $\varepsilon < 0$ and in case of untargeted attack, y are incorrect targets and $\varepsilon > 0$.

In [9], it was discovered that adding the momentum term into the iterative process for attacks lead to more stable optimization trajectory. It is determined that adversarial examples generated with the iterative method with the use of momentum are more suitable for white-box attacks than the ones generated without the use of momentum.

In [10], the question of the robustness of adversarial examples under real-world image transformations was studied, and Expectation Over Transformation (EOT) algorithm was proposed. This algorithm helps to generate an adversarial example taking into account a set of transformations which usually spoils the transferability of an image to the real-world.

Given the distance function $\rho : X \times X \rightarrow \mathbb{R}$, distance bound ε , set of transformations T and y_{adv} as desired adversarial

classes corresponding to x , EOT approach may be formulated as the following optimization problem:

$$X_{\text{adv}} = \arg \max_z \mathbb{E}_{t \in T} \left[\mathbb{P}(y_{\text{adv}} | t(x+z)) \right], \quad (2)$$

such that $\mathbb{E}_{t \in T} [\rho(x+z, x)] < \varepsilon$ and $x+z \in [0, 1]^d$. In this formulation $z \in [0, 1]^d$ is an additive noise. This approach helps to imitate such transformations as camera noise or viewpoint shifts.

III. METHODOLOGY

It should be mentioned that usually face recognition process in the real world consists of two sequential parts. Firstly, the image of a person is passed to face detection neural network, which obtains a bounding box with a face. Secondly, the obtained bounding box is passed to the classifying neural network, LResNet100E-Ir in our case. Nowadays, there is a variety of face detection neural networks. In our experiments, we use the multi-task cascaded CNN based framework (MTCCN), proposed in [11].

In this paper, we propose a pipeline to generate a broad class of adversarial patches, which can fool white-box models and should be fitted for a given person. In this section, we describe the proposed pipeline and give a detailed explanation of each stage.

The attacking strategy may be explained as a sequence of the following steps:

- choose desired patch location and its shape,
- print patch with a chessboard pattern and apply it on a face or wearable attribute,
- take a photo of your face and mark corners of the chessboard on a photo,
- repeat steps mentioned before to get photos with different head rotations,
- project adversarial patch on the chessboard pattern,
- get facial landmarks using face detector (MTCNN),
- warp image to standard 112×112 ArcFace input using similarity transformation and obtained landmarks,
- solve the corresponding optimization problem.

A. Projective transformation and chessboard pattern

1) *Non-linear 2d transformation*: To approximate non-linear 2d transformation, we apply projective transformation separately to each cell of the chessboard grid. Regular grid $G = \{G_i\}$ of pixels $G_i = (x_i^t, y_i^t)$ forming output map. (x_i^t, y_i^t) denotes chessboard sticker's coordinates on the photo. By four corner pixels of j 'th chessboard cell and corresponding four marked pixels of projected cell we can obtain M_{θ_j} matrix of spatial transformation, so sampling points of j 'th cell:

$$\begin{pmatrix} x_{i_j}^s \\ y_{i_j}^s \end{pmatrix} = \mathcal{T}_{\theta_j}(G_{i_j}) = M_{\theta_j} \begin{pmatrix} x_i^t \\ y_i^t \\ 1 \end{pmatrix}, \quad (3)$$

where (x_i^s, y_i^s) are the source coordinates in the input map, they define the spatial location in the adversarial patch where sampling kernel is applied to get the output map. Given the expression above, the target grid for the adversarial patch is obtained as a union of target coordinates of each cell:

$$\begin{pmatrix} x_i^s \\ y_i^s \end{pmatrix} = \{\mathcal{T}_{\theta_j}(G_{i_j})\} \quad (4)$$

To speed up the training procedure, we should precalculate and store sampling points. Next step is differentiable image sampling. In more details, this procedure described in [12]. Scheme of transformations is illustrated in **Fig. 1**.

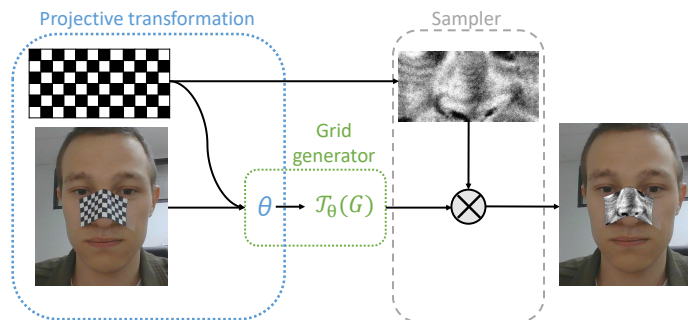


Fig. 1. Procedure of adversarial sticker application.

2) *Linear 2d transformation*: In case when the application of adversarial patch could be simulated by a linear transformation (e.g. in case of eyeglasses), we repeat steps mentioned in the previous section just for one cell.

B. Similarity transform and face detector

MTCNN face detection system finds facial landmarks, which are then used to obtain a matrix of similarity transformation A_ϕ and sampling points $\mathcal{T}_\phi(G')$ for this warp. The main reason why it is necessary to apply sticker before similarity transformation is image sampling. Sampling kernel takes values of neighbouring pixels and copies one of them or return averaged value. Application patch after similarity transformation tends to inconsistent values of pixels on the boundaries of the patch.

C. Loss function and training procedure

The goal of the attack is to generate adversarial patch which can not only mislead face recognition system but also should be good-looking. We formulate the second objective by restricting neighbouring pixels to have a similar colour. We obtain the adversarial sticker minimizing the following objective function:

$$\mathcal{L}(X, p) = \mathcal{L}_{\text{adv}}(X, p) + \tau TV(p), \quad p \in [0, 1]^{m \times n} - \text{patch}, \quad (5)$$

where X is a batch of photos of the attacker with different shooting conditions (i.e. set of transforms T described above) and total variation loss, or TV loss, preserve high perceptual quality and ensures that the difference between neighbouring pixels is imperceptible [13], as far as photos taken from the camera do not have tremendous shift in values of neighbouring pixels.

TV loss is the function of pixel values of a patch p and is defined as follows:

$$TV(p) = \sum_{i,j} \sqrt{(p_{i,j} - p_{i,j+1})^2 + (p_{i,j} - p_{i+1,j})^2} \quad (6)$$

\mathcal{L}_{adv} is the cosine similarity loss that minimizes the similarity between an embedding of the photo with patch and initial embedding of the person. This term depends on type of attack. In case of untargeted attack:

$$\mathcal{L}_{\text{adv}}(X, p) = \mathbb{E}_{t \in T, x \in X} [-\cos(e_{x_t}, e)], \quad (7)$$

where x is a photo from batch X corresponding to some shooting condition, e is ground truth embedding of attacker.

In case of targeted attack:

$$\mathcal{L}_{\text{adv}}(X, p) = \mathbb{E}_{t \in T, x \in X} [\cos(e_{x_t}, e_{x'})], \quad (8)$$

where $e_{x'}$ is an embedding corresponding to desired person or to the closest class except for ground truth class.

Here e_{x_t} is an embedding corresponding to the photo x_t of attacker with applied patch. Procedure of patch application may be denoted in a following way. Given a sampler function A , photo x and corresponding grid $\mathcal{T}_\theta(G)$, x_t denotes image with applied patch:

$$x_t = A(x, t(p), \mathcal{T}_\theta(G)). \quad (9)$$

It should be mentioned that all the operations described above are differentiable. According to [9], gradients through this iterations $\nabla_p \mathcal{L}(X, p)$ are accumulated. Given decay factor μ ,

$$g := \mu g + \frac{\nabla_p \mathcal{L}(X, p)}{\|\nabla_p \mathcal{L}(X, p)\|_1} \quad (10)$$

and patch p is updated as follows:

$$p := p - \varepsilon \text{sign}(g). \quad (11)$$

The whole attacking pipeline is presented in **Fig 2**.

IV. EXPERIMENTS

In this section, we describe settings of experiments and its technical details.




Fig. 2. Attacking pipeline

A. Data preprocessing

As a part of the experimental setup, we used pretrained LResNet100E-IR to collect embeddings of different people. For this purpose, we use first 200 classes from CASIA-WebFace [14] dataset and photos of the 1st and the 2nd authors as images of attackers.

B. Setup of experiments

It should be mentioned that solving the optimization problem described in the previous section implies patch optimization in the whole RGB colour space. However, colour triplets of the patch in digital space differ significantly from the ones of the printed patch because of the printer's narrow colour spectrum. More than that, triplets change one more time when the printed patch is photographed on camera. Due to these issues, we decided to use not a colour but a grayscale patch.

It was discovered that iterative FGSM method with momentum works good in the digital domain, but gradient step size ε , decay factor μ and weight τ of TV loss affect colours transition to physical domain dramatically. It was empirically estimated that values $\varepsilon = \frac{1}{16}$, $\mu = 0.9$ and $\tau = 1e - 3$ during the training leads to the best results on test in the real world.

In addition, we want to mention that initial similarity with ground truth class $\cos(e_{x_t}, e)$ is between 0.65 and 0.70 for all training photos in all three scenarios, whereas initial similarity with desired class $\cos(e_{x_t}, e_{x'})$ in targeted attacks is always below 0.30.

C. Adversarial sticker localization

In this work, we tested the method proposed above to generate patches for three different areas on the face. For the diversity, one of the generated patches is an eyeglass; two others are stickers on nose and forehead.

1) *Attacking eyeglasses:* In this case, we use a patch of the form of eyeglasses rim (front view), where the frame and bridge of eyeglasses are used for adversarial pattern application. It turned out that the size of these parts is crucial for the success of the attack. Thus, we decided to use an eyeglasses model with a big frame and of size 16.5×6.4 cm. Example of the adversarial patch of this form is presented in Fig. 3 and Fig. 4.



Fig. 3. Adversarial eyeglasses

In this scenario, 14 photos were used for training, 3 photos were used for validation in the digital domain, and 5 photos were used for the test in the physical domain.

2) *Attacking forehead:* In this case, we used a rectangular patch of size 14×5 cells; the width of one cell is 1.0 cm. It turned out that the closer patch to eyebrows the better results in a targeted attack. In this scenario, 4 photos were used for



Fig. 4. Example of adversarial eyeglasses

training, 3 photos were used for validation in digital domain and 3 photos were used for the test in the physical domain. The patch of this type is illustrated in Fig. 5 and in Fig. 6.



Fig. 5. Adversarial forehead



Fig. 6. Example of adversarial sticker on forehead

3) *Attacking nose:* In this case, we used a rectangular patch of size 12×6 cells; the width of one cell is 0.7 cm. An example of adversarial sticker is depicted in Fig. 7 and Fig. 8. We varied length in cells of the patch. During experiments, we found out that the height of the patch less than 6 cell is not enough for the successful attack even in the digital domain. but it was found out that increasing of size of patch is not enough for transferability of attack to the physical domain.

D. Results

In all three above scenarios, training procedure was conducted for 10000 iterations of the method described before. If on some iteration $\mathcal{L}_{adv}(X, p) < 0$, training stops and an obtained patch is printed and tested in the physical world.

Numerical results of the attacks for the first and the second attacking scenarios are presented in Table I. There we report the mean similarities \pm standard errors of the mean on training photos x_{train} , on validation photos in digital domain x_{val} and on test photos in physical domain x_{test} with both ground truth class and desired class.




Fig. 7. Adversarial nose



Fig. 8. Example of adversarial sticker on nose

E. Discussion

Numerical experiments showed that it is possible to attack ArcFace in the real world with the use of adversarial stickers placed on eyeglasses or forehead. Although similarity with ground truth class $\cos(e_{x_{test}}, e)$ is just slightly below the similarity with desired class $\cos(e_{x_{test}}, e_{x'})$, it is noticeable that neural network can not recognize the attacker as ground truth class. It should be said that the improvement of stability of the proposed technique to different illumination conditions may increase the success of the attack.

In addition, it may be mentioned that one of the interesting outcomes of attack is appearing of sketches of parts of the face on adversarial patches located in corresponding areas. More than that, it is discovered that the position of a patch, as well as its size, dramatically affects the success of attack in the physical domain. In general, the bigger the patch and the closer it to eyes, the better the results of attack in the real world.

V. CONCLUSION AND FUTURE WORK

We proposed a simple method of creating adversarial patches that can be used to attack state-of-the-art face recognition systems. Our approach was tested in targeted and untargeted attacks on FaceID model LResNet100E-IR, ArcFace@ms1m-refine-v2. In this paper, it was shown that a simple and easily reproducible attacking technique leads to incorrect work of that system not only in the digital domain but also in the physical world. Experiments show that further security research should be conducted to improve the robustness of such networks to attacks of this type. In the future, we focus on construction and utilization of a defense

TABLE I
NUMERICAL RESULTS OF EXPERIMENTS

Patch	Type of attack	$\cos(e_{x_{train}}, e)$	$\cos(e_{x_{train}}, e_{x'})$	$\cos(e_{x_{val}}, e)$	$\cos(e_{x_{val}}, e_{x'})$	$\cos(e_{x_{test}}, e)$	$\cos(e_{x_{test}}, e_{x'})$
Eyeglasses	Targeted	0.041 ± 0.052	0.648 ± 0.020	0.317 ± 0.004	0.451 ± 0.021	0.305 ± 0.024	0.363 ± 0.024
Forehead	Targeted	-0.053 ± 0.009	0.221 ± 0.011	0.273 ± 0.007	0.421 ± 0.025	0.323 ± 0.035	0.391 ± 0.021

technique to make a network invulnerable to such attacks, especially in the physical world.

REFERENCES

- [1] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4690–4699.
- [2] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.
- [3] T. B. Brown, D. Mané, A. Roy, M. Abadi, and J. Gilmer, "Adversarial patch," *arXiv preprint arXiv:1712.09665*, 2017.
- [4] B. Heisele and T. Koshizen, "Components for face recognition," in *Sixth IEEE International Conference on Automatic Face and Gesture Recognition, 2004. Proceedings.* IEEE, 2004, pp. 153–158.
- [5] M. Sharif, S. Bhagavatula, L. Bauer, and M. K. Reiter, "Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security.* ACM, 2016, pp. 1528–1540.
- [6] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," *arXiv preprint arXiv:1607.02533*, 2016.
- [7] S. Qiu, Q. Liu, S. Zhou, and C. Wu, "Review of artificial intelligence adversarial attack and defense technologies," *Applied Sciences*, vol. 9, no. 5, p. 909, 2019.
- [8] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *CoRR*, vol. abs/1412.6572, 2014.
- [9] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, and J. Li, "Boosting adversarial attacks with momentum," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 9185–9193.
- [10] A. Athalye, L. Engstrom, A. Ilyas, and K. Kwok, "Synthesizing robust adversarial examples," in *ICML*, 2017.
- [11] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, 2016.
- [12] M. Jaderberg, K. Simonyan, A. Zisserman *et al.*, "Spatial transformer networks," in *Advances in neural information processing systems*, 2015, pp. 2017–2025.
- [13] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: nonlinear phenomena*, vol. 60, no. 1-4, pp. 259–268, 1992.
- [14] D. Yi, Z. Lei, S. Liao, and S. Z. Li, "Learning face representation from scratch," *arXiv preprint arXiv:1411.7923*, 2014.