THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Transformation on Computer-Generated Facial Image to Avoid Detection by Spoofing Detector

Citation for published version:

H. Nguyen, H. T. Tieu, N-D, Nguyen-Son, H-Q, Yamagishi, J & Echizen, I 2018, Transformation on Computer-Generated Facial Image to Avoid Detection by Spoofing Detector. in *IEEE International Conference on Multimedia and Expo (ICME) 2018*. Institute of Electrical and Electronics Engineers (IEEE), San Diego, USA, IEEE International Conference on Multimedia and Expo 2018, San Diego, California, United States, 5/07/18. <https://doi.org/10.1109/ICME.2018.8486579>

Digital Object Identifier (DOI):

[10.1109/ICME.2018.8486579](https://doi.org/10.1109/ICME.2018.8486579)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

IEEE International Conference on Multimedia and Expo (ICME) 2018

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

TRANSFORMATION ON COMPUTER-GENERATED FACIAL IMAGE TO AVOID DETECTION BY SPOOFING DETECTOR

Huy H. Nguyen¹, Ngoc-Dung T. Tieu¹, Hoang-Quoc Nguyen-Son², Junichi Yamagishi^{1,2,3}, Isao Echizen^{1,2}

¹Graduate University for Advanced Studies, Kanagawa, Japan

²National Institute of Informatics, Tokyo, Japan

³The University of Edinburgh, Edinburgh, UK

{nhhuy, dungtieu, nshquoc, jyamagishi, iechizen}@nii.ac.jp

ABSTRACT

Making computer-generated (CG) images more difficult to detect is an interesting problem in computer graphics and security. While most approaches focus on the image rendering phase, this paper presents a method based on increasing the naturalness of CG facial images from the perspective of spoofing detectors. The proposed method is implemented using a convolutional neural network (CNN) comprising two autoencoders and a transformer and is trained using a black-box discriminator without gradient information. Over 50% of the transformed CG images were not detected by three state-of-the-art spoofing detectors. This capability raises an alarm regarding the reliability of facial authentication systems, which are becoming widely used in daily life.

Index Terms— deep convolutional neural network, autoencoder, presentation attack, computer-generated image, spoofing detection

1. INTRODUCTION

A presentation attack is commonly used to bypass authentication systems using biometrics information (face, fingerprint, iris, and/or voice). Integration of a spoofing detector into the system before the authentication phase is one approach to preventing such attacks. A good candidate for this is liveness detection, which generally uses a challenge-response protocol in which the user is asked to perform an action such as blinking, smiling, or moving the lips. However, recent work has shown that it is possible to avoid liveness detection by, for example, using real-time face capture and reenactment [1]. It has thus become necessary to develop and implement natural-CG image/video discriminators.

Forensic research on discriminating between CG and natural images has focused on both images in general and facial images. As an example of the former, Wu et al. extracted statistical features from histograms of differential images [2]. Although this approach was proposed several years ago, our evaluation demonstrated that it has fast feature extraction and good performance. Therefore, we used it as the basis of the discriminator used to train our CNN. In 2017, Peng et al. [3] reported a method based on multi-fractal and regression analysis. As an example of focusing on facial images, the recent work of Nguyen et al. [4] focused

Fig. 1: Scenario of proposed presentation attack method.

on facial smoothness as represented by edges and local entropy of the skin areas.

In the research reported here, we developed a method for avoiding detection by spoofing detectors like those natural-CG image discriminators mentioned above. It works by transforming CG images input to facial authentication systems to make them appear more natural. Our work is motivated by the idea of “adversarial machine learning” of Huang et al. for attacking machine learning based systems [5]. Although we did not target a specific system, we used a spoofing detection algorithm as the basis of the discriminator. The attack scenario is illustrated in Fig. 1. Given two sets of data (a set of natural images and a set of CG ones which are not necessarily corresponding person-to-person or pose-to-pose), a system using the proposed method implemented as a CNN transforms the CG input images in an attempt to make them indistinguishable from the natural counterparts.

Unlike the original generative adversarial network (GAN) [6], the discriminator used to train our CNN was pre-trained, kept fixed during training, and could not perform back propagation. Moreover, the generated images retained important information from the input images, such as the persons identity, the facial expression, or the lips’ shapes, which are

Fig. 2: H-Net architecture.

hard to control when using the original GAN. When designing our CNN we focused on minimizing the number of parameters so that it can work without consuming a large amount of GPU memory. We also considered the “training with small dataset problem,” which is often faced by attackers when collecting data.

Our approach is different from those of other computer graphic engines, which mainly focus on the rendering phase. It also differs from the style transfer problem, which mainly focuses on applying the “look” or textures of one image (the style) to another one (the content). In our case, the style is not just a picture or an artistic style, so it is hard to define.

Our contributions here are threefold:

- Presentation of a CNN comprising two autoencoders and a transformer net that increases the difficulty of detecting computer-generated facial images.
- Suggestion of a method for dealing with back propagation when training using an external black-box discriminator that does not have gradient information.
- Raising of an alarm about the robustness of facial authentication systems, which are being implemented in many mobile devices and have become a tempting target for attacker.

2. MODEL ARCHITECTURE

2.1. Overview

It is very hard to explicitly point out what makes natural images look “natural” and CG ones look “unnatural.” We humans have been “trained” by our encounters with many natural things and scenes since we were born, so we can intuitively distinguish which images were produced by computer. Some forensic researchers have tried describing these intuitive feelings in terms of specific properties [3, 4]. Others, such as Wu et al. [2], have tried using statistical methods to distinguish natural images from CG images. However, these approaches are problematic, require lengthy experiments and do not sufficiently describe the essence of the two kinds of images.

Fig. 3: Design of encoder net (left) and decoder net (right).

To tackle this feature-finding problem, we use two autoencoders [7] with the same design but different weights ($\{E_{\theta_{EN}}, D_{\theta_{DN}}\}$ and $\{E_{\theta_{CG}}, D_{\theta_{DCG}}\}$) to automatically extract latent features from natural image I_N and CG image I_{CG} . They are trained to learn representations of the inputs and to regenerate them with the same properties. To transform the latent feature space of CG images into those of natural ones, we use a transformer net T_{θ_T} with five bottleneck residual blocks [8]. This block design reduces the number of parameters for the network. Old information given by the skip connection combines with new one created by the layers in a residual block are suitable for transformation. Since our network is shaped like an “H,” as illustrated in Fig. 2, we call it “H-Net.” In the evaluation phase, the CG encoder $E_{\theta_{CG}}$, the transformer T_{θ_T} , and the natural decoder $D_{\theta_{DN}}$ transform CG input I_{CG} into a natural-looking image O_T that is similar to natural image I_N from the perspective of a natural–CG image discriminator.

The encoder has four convolutional layers with 3×3 kernels and a stride of 1. Batch normalization [9] is used to deal with high learning rates and the less careful initialization problem. After batch normalization layers are exponential linear unit (ELU) layers, which give better performance than traditional rectified linear units (ReLU), leaky ReLUs (LReLU), and parametrized ReLUs (PReLU) [10]. To reduce the spatial size of the representation, we use 2×2 average pooling layers, which have recently come to be used instead of max pooling ones. In the decoder, the convolutional layers are replaced with their counterparts: transposed convolutional ones and two sub-pixel convolutional ones, which have been reported to give better upscaling results than previous approaches[11]. The detailed designs of the encoder

and decoder are shown in Fig 3.

2.2. Loss functions

Loss functions are used to optimize the model in the training phase. Pixel-wise loss functions such as mean squared error (MSE) ones are widely used but have limited capabilities for images when measuring perceptual quality [12]. To overcome this limitation, Dosovitskiy and Brox [13] used Euclidean-based loss in feature space in combination with adversarial loss. Another approach is to use a perceptual loss function based on the Euclidean distance between feature maps extracted from a VGG network proposed by the Visual Geometry Group at the University of Oxford [14]. Ledig et al. combined both VGG loss and adversarial loss in their super-resolution generative adversarial network [15].

We use four loss functions:

1. MSE loss L_{MSE} is only used to calculate the loss between two feature maps, including high-level presentations encoded by the encoders and others extracted from the VGG-19 network.
2. VGG loss L_{VGG} , or perceptual loss, is the MSE loss of features maps (of I_1 and I_2) extracted using the VGG-19 network: $L_{VGG} = L_{MSE}(VGG(I_1), VGG(I_2))$. VGG loss represents the perceptual quality of generated images.
3. Adversarial loss L_{Adv} is the binary cross entropy loss between two labels: one is from the discriminator and the other is the destination label. Details of the discriminator are discussed in section 2.3.
4. Total loss L_{Tot} is the combination of L_{VGG} and L_{Adv} , formulated as equation 1. It represents the quality of images generated by natural decoder $D_{\theta_{EN}}$. In our experiments, we set the hyper-parameter α to 5×10^{-3} . Gradient back propagation is discussed in section 2.3.

$$L_{Tot} = (1 - \alpha)L_{VGG} + \alpha L_{Adv} \quad (1)$$

2.3. Discriminator

Unlike the original GAN proposed by Goodfellow et al. [6], we use a pre-trained discriminator for training our H-Net. It can be any black-box spoofing detector. Since it is unlikely for attackers in the wild to have full access to the spoofing detectors in authentication systems, it is more realistic to assume that the spoofing detector used as the discriminator used for training is a “black-box” rather than a differentiable white-box discriminator. In our experiments, the use of a “traditional” discriminator in the GAN, such as the one used by Ledig et al. [15], results in very poor performance. This is because performing only convolution is not enough for extracting the features needed for distinguishing natural images from CG images. Hence, we borrow the method of Wu et al., which focuses on statistical features from histograms of differential images [2]. We use a neural net instead of Fisher’s linear discriminant analysis (FLDA) as Wu et al. did because

we can easily run the discriminator in parallel with H-Net in the training phase without using an additional framework.

We need to compute the gradient of total loss L_{Tot} in equation 1 with respect to the parameters in the generator. However, since the discriminator is a black-box, the adversarial loss L_{Adv} is not differentiable. One possible approach is to approximate it on the basis of the gradient of the VGG loss L_{VGG} only. However, the adversarial loss makes no contribution to the gradient. Therefore, we approximate the total gradient as follows:

$$grad_{L_{Tot}} = \frac{(1 - \alpha)L_{VGG} + \alpha L_{Adv}}{L_{VGG}} grad_{L_{VGG}} \quad (2)$$

2.4. Training and evaluation

H-Net is trained using two sets of images: natural images and CG ones. The two sets do not have to be pairwise correlated, which means that they may contain images of different people with different poses or facial expressions. This loose condition comes from reality; i.e., it is very hard for an attacker to find a large number of natural–CG image pairs with similar perceptual content.

In training, four optimization processes are performed sequentially:

Step 1: Training the natural autoencoder $\{E_{\theta_{EN}}, D_{\theta_{DN}}\}$ with natural image I_N by minimizing equation 3.

$$\argmin_{\theta_{EN}, \theta_{DN}} L_{Tot}[D_{\theta_{DN}}(E_{\theta_{EN}}(I_N)), I_N] \quad (3)$$

The output image $O_N = D_{\theta_{DN}}(E_{\theta_{EN}}(I_N))$ must satisfy two conditions: (1) have the same perceptual content as input I_N and (2) be classified as a natural image by the discriminator.

Step 2: Training the CG autoencoder $\{E_{\theta_{ECG}}, D_{\theta_{DCG}}\}$ with CG image I_{CG} by minimizing equation 4.

$$\argmin_{\theta_{ECG}, \theta_{DCG}} L_{VGG}[D_{\theta_{DCG}}(E_{\theta_{ECG}}(I_{CG})), I_{CG}] \quad (4)$$

In this step, the discriminator loss is not necessary. The only requirement is that output image O_{CG} be perceptually similar to I_{CG} .

Step 3: Training the transformer net T_{θ_T} to convert latent features encoded by CG encoder $E_{\theta_{ECG}}$ into ones that have the same distribution of features as those extracted by the natural encoder $E_{\theta_{EN}}$. As mentioned above, due to the lack of pairwise correlated I_N and I_{CG} pairs, it is very hard to minimize the loss between $E_{\theta_{EN}}(I_N)$ and $T_{\theta_T}(E_{\theta_{ECG}}(I_{CG}))$. An acceptable solution is feeding I_N into both networks and minimizing the features they encode, formulated in equation 5.

$$\argmin_{\theta_T} L_{MSE}[E_{\theta_{EN}}(I_N), T_{\theta_T}(E_{\theta_{ECG}}(I_N))] \quad (5)$$

Step 4: Training the CG transformation path, which includes CG encoder $E_{\theta_{ECG}}$, transformer T_{θ_T} , and natural decoder $D_{\theta_{DN}}$. As in the first step, total loss must be used to ensure that transformed output $O_T = D_{\theta_{DN}}(T_{\theta_T}(E_{\theta_{ECG}}(I_{CG})))$ is not classified as a CG image. The lack of pairwise correlated I_N and

I_{CG} pairs in Step 3 appears here as well, so we need to replace I_N with I_{CG} in equation 6.

$$\underset{\theta_{E_{CG}}, \theta_T, \theta_{D_N}}{\operatorname{argmin}} L_{Tot}[D_{\theta_{D_N}}(T_{\theta_T}(E_{\theta_{E_{CG}}}(I_{CG}))), I_{CG}] \quad (6)$$

Doing this could also bring back some “balance” from the effect of I_N on CG encoder $E_{\theta_{E_{CG}}}$ in step 3.

When we compute the transformed output from the CG input in the evaluation phase, we use a step similar to step 4 in the training phase:

$$O_T = D_{\theta_{D_N}}(T_{\theta_T}(E_{\theta_{E_{CG}}}(I_{CG}))) \quad (7)$$

We do not use the $\tanh$ function to scale the output images because of the resulting low contrast. To avoid extreme values with resulting abnormal color spots, the pixel intensities are restricted to the range $[-1.8, 1.8]$ before being denormalized to $[0, 255]$.

3. EVALUATION

In facial image forensic research, there was no standardized dataset used for mutual comparison. Each research group tended to have its own datasets. Therefore, we had to combine pieces from five different sources to create three datasets used for evaluation (see Table 1). All images were resized to 256×256 pixel resolution. We tested H-Net on two scenarios: (1) the attacker knows the dataset used for training the spoofing detector, and (2) the attacker has no knowledge of the training dataset. Three spoofing detectors were used: Wu et al.’s [2], Peng et al.’s [3], and Nguyen et al.’s [4]. Note that only Wu et al.’s spoofing detector was used as a discriminator for training H-Net; the other spoofing detectors were unseen by the attacker. Moreover, in Wu et al.’s one, we used FLDA for classification as in the authors’ report. We compared both the accuracies and detection rates of the three spoofing detectors. Let n_{TP} and n_{TN} be respectively the number of images correctly classified as CG or natural, n_{FN} be the number of CG or transformed images misclassified as natural ones, and N be the total number of images. Accuracy is defined as sum of n_{TP} and n_{TN} over N : $\frac{n_{TP}+n_{TN}}{N}$. The detection rate represents the ability of the spoofing detector to detect positive items: $\frac{n_{TP}}{n_{TP}+n_{FN}}$.

3.1. Scenario 1: Attacker knows training dataset of spoofing detector

We trained both spoofing detectors and H-Net on dataset 1. We then evaluated them on all datasets to see if the images transformed by H-Net could avoid detection by these pre-trained spoofing detectors. Comparisons of sample images before and after transformation are shown in Figs. 4 and 5. As shown in Table 2, the detection rates significantly decreased when the CG images were transformed, especially for Wu et al.’s and Peng et al.’s methods.

Although the method of Nguyen et al. had the lowest performance on dataset 1, its detection rate after transformation was the highest (nearly 50%). On datasets 2 and 3, its performance before transformation was very poor; it was better after transformation. Our analysis shows that this method was over-fitted

Fig. 4: Original images (top) and transformed ones (bottom) from dataset 1. Color and brightness of images in second row were normalized by H-Net as learned from training data. First two images demonstrate good transformation in perception. Contrast of third image was improved. In last image, skin color was whitened a bit undesirably.

Fig. 5: Original images (top) and transformed ones (bottom) from dataset 2 (left) and dataset 3 (right). Brightness of first two image was unexpectedly reduced due to bright background. Although grayscale images were given skin-like color, they can easily be converted back into grayscale.

©Copyright University of Basel. ©Copyright 2003-2005 Massachusetts Institute of Technology. All Right Reserved.

for dataset 1 which had high-quality CG images. It had a tendency to classify fine-texture images as CG ones. On datasets 2 and 3, which did not have fine-texture CG images, it classified almost of the images as natural ones. Because of the spoofing detector’s preset threshold, some transformed images had good enough texture to be classified as CG ones, which increased the detection rate. Therefore, if an attacker tries to avoid detection by this spoofing detector, he or she may be successful the first time without the help of H-Net.

3.2. Scenario 2: Attacker does not know training dataset of spoofing detector

Unlike in the first scenario, we trained the H-Net and the spoofing detectors on difference datasets (switched between dataset 1 and

Table 1: Three datasets used in evaluation.

No.	Components	Size	Description
1	Dang-Nguyen et al. [16]	CG: 240 Nat: 240	40 very realistic CG images collected from Web plus 200 good quality CG images extracted from PES 2012 soccer game 240 natural images retrieved from Internet
2	Basel (CG) [17] Caltech99 (Nat) [18]	CG: 270 Nat: 270	3D face scans and rendered images from Basel Face Model Natural images from Caltech Faces 1999 dataset
3	MIT (CG - Grayscale) [19] MS-Celeb-1M (Nat) [20]	CG: 3236 Nat: 3236	CG images extracted from MIT CBCL dataset Natural images selected from MS-Celeb-1M cropped version

Table 2: Scenario 1 - Accuracies (%) and detection rates (%) of three spoofing detectors on three datasets before and after performing transformation on CG parts.

Spoofing detectors	Dataset 1				Dataset 2				Dataset 3			
	Accuracy		Detection rate		Accuracy		Detection rate		Accuracy		Detection rate	
	Before	After↓	Before	After↓	Before	After↓	Before	After↓	Before	After↓	Before	After↓
Wu et al. [2]	92.71	48.33	93.33	0.00	83.89	55.37	93.33	36.30	64.65	56.83	35.51	19.90
Peng et al. [3]	90.63	52.71	92.08	16.25	59.26	19.26	97.41	17.41	50.20	13.58	100.00	26.79
Nguyen et al. [4]	83.54	64.79	87.08	48.33	17.78	32.78	0.00	30.00	32.31	63.23	0.49	62.11

2), and evaluated them on dataset 3.

3.2.1. Scenario 2.1: H-Net was trained on dataset 1, spoofing detectors were trained on dataset 2

The evaluation results on dataset 3 are shown in Table 3. The transformed images again significantly reduced the detection rates of all spoofing detectors, especially those of Wu et al. and Peng et al., which were nearly 0%. Nguyen et al.'s method had a detection rate of around 50%, down from nearly 100%, meaning that the attacker had a 50–50 chance of avoiding detection by this spoofing detector. In this case, this spoofing detector learned that low-texture images had a high probability of being CG ones, which was opposite to its knowledge in scenario 1. Therefore, after transformation on dataset 3, the textures of the CG images were improved so that the images would likely be classified as natural ones. This also clarifies the contrasting changes in the detection rate in the first scenario on dataset 1 vs. datasets 2 and 3.

Table 3: Scenario 2.1 - Evaluation results on dataset 3.

Spoofing detectors	Accuracy		Detection rate	
	Before	After	Before	After
Wu et al. [2]	56.38	6.46	100.00	0.19
Peng et al. [3]	92.32	42.57	100.00	0.49
Nguyen et al. [4]	96.72	71.54	99.20	48.89

3.2.2. Scenario 2.2: H-Net was trained on dataset 2, spoofing detectors were trained on dataset 1

The evaluation results on dataset 3 are shown in Table 4. Before transformation, Peng et al.'s method seemed to classify all input as CG. After transformation, its decision was changed that all transformed CG images were classified as natural ones, therefore

both the accuracy and the detection rate were around 0. Nguyen et al.'s method had the same behavior as in the first scenario as expected. The performance of Wu et al.'s method was increased after transformation, which was different from the two above scenarios. A possible explanation for this phenomenon is that the training dataset used for H-Net is small and monotonous. As the result, H-Net did not have enough knowledge about other kinds of CG images.

Table 4: Scenario 2.2 - Evaluation results on dataset 3.

Spoofing detectors	Accuracy		Detection rate	
	Before	After	Before	After
Wu et al. [2]	64.65	82.73	35.51	71.66
Peng et al. [3]	50.20	0.20	100.00	0.00
Nguyen et al. [4]	32.31	41.27	0.49	18.17

4. CONCLUSION AND FUTURE WORK

The performances of both H-Net and detectors are depended on the quality of training datasets. However, in most cases, over 50% of the CG images transformed using our H-Net avoided detection by three state-of-the-art spoofing detectors. Since the facial features were preserved, facial recognition was unaffected. This means that the network can be trained using a black-box discriminator that cannot perform back propagation. However, H-Net has some limitations, especially when up to 50% transformed images are still separable from natural ones. Future work will mainly focus on this limitation as well as finding better datasets. We will also evaluate the local substitute method to perform black-box attacks [21]. Other issues to be addressed are solving the black skin-color problem, dealing with larger images, and reducing network size to enable it to work smoothly with video frames in real time.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Numbers JP16H06302 and 15H01686.

We are grateful to the Massachusetts Institute of Technology and to the Center for Biological and Computational Learning for providing databases of facial images. We thank Prof. Dr. Thomas Vetter, Department of Computer Science at the University of Basel, for the 3D face scans and renderings dataset. We also thank the Computer Vision Department of Caltech University, Microsoft Corporation, Dr. Duc-Tien Dang-Nguyen and his colleagues for providing invaluable facial datasets for conducting our experiments.

5. REFERENCES

- [1] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner, “Face2Face: Real-time face capture and reenactment of RGB videos,” in *CVPR*, 2016, pp. 2387–2395.
- [2] Ruoyu Wu, Xiaolong Li, and Bin Yang, “Identifying computer generated graphics via histogram features,” in *ICIP*. IEEE, 2011, pp. 1933–1936.
- [3] Fei Peng, Die-lan Zhou, Min Long, and Xing-ming Sun, “Discrimination of natural images and computer generated graphics based on multi-fractal and regression analysis,” *AEU-International Journal of Electronics and Communications*, vol. 71, pp. 72–81, 2017.
- [4] Huy H Nguyen, Hoang-Quoc Nguyen-Son, Thuc D Nguyen, and Isao Echizen, “Discriminating between computer-generated facial images and natural ones using smoothness property and local entropy,” in *IWDW*. Springer, 2015, pp. 39–50.
- [5] Ling Huang, Anthony D Joseph, Blaine Nelson, Benjamin IP Rubinstein, and JD Tygar, “Adversarial machine learning,” in *ACM Workshop on Security and Artificial Intelligence*. ACM, 2011, pp. 43–58.
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in *NIPS*, 2014, pp. 2672–2680.
- [7] Yoshua Bengio et al., “Learning deep architectures for AI,” *Foundations and trends® in Machine Learning*, vol. 2, no. 1, pp. 1–127, 2009.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *CVPR*, 2016, pp. 770–778.
- [9] Sergey Ioffe and Christian Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *ICML*, 2015, pp. 448–456.
- [10] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter, “Fast and accurate deep network learning by exponential linear units (ELUs),” *ICLR*, 2016.
- [11] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang, “Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network,” in *CVPR*, 2016, pp. 1874–1883.
- [12] Michael Mathieu, Camille Couprie, and Yann LeCun, “Deep multi-scale video prediction beyond mean square error,” *ICLR*, 2016.
- [13] Alexey Dosovitskiy and Thomas Brox, “Generating images with perceptual similarity metrics based on deep networks,” in *NIPS*, 2016, pp. 658–666.
- [14] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” *ICLR*, 2015.
- [15] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al., “Photo-realistic single image super-resolution using a generative adversarial network,” *Computing Research Repository (CoRR)*, 2016.
- [16] Duc-Tien Dang-Nguyen, Giulia Boato, and Francesco GB De Natale, “Discrimination between computer generated and natural human faces based on asymmetry information,” in *EUSIPCO*. IEEE, 2012, pp. 1234–1238.
- [17] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter, “A 3D face model for pose and illumination invariant face recognition,” in *Advanced Video and Signal Based Surveillance*. IEEE, 2009, pp. 296–301.
- [18] M Weber and M Weber, “Caltech frontal face database,” *California Institute of Technology*, 1999.
- [19] Benjamin Weyrauch, Bernd Heisele, Jennifer Huang, and Volker Blanz, “Component-based face recognition with 3D morphable models,” in *CVPR Workshop*. IEEE, 2004, pp. 85–85.
- [20] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao, “MS-Celeb-1M: A dataset and benchmark for large scale face recognition,” in *ECCV*. Springer, 2016, pp. 87–102.
- [21] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami, “Practical black-box attacks against machine learning,” in *AsiaCCS*. ACM, 2017, pp. 506–519.