

# Real-world Attack on MTCNN Face Detection System

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

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

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

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

**Abstract**—Recent studies proved that deep learning approaches achieve remarkable results on face detection task. On the other hand, the advances gave rise to a new problem associated with the security of the deep convolutional neural network models unveiling potential risks of DCNNs based applications. Even minor input changes in the digital domain can result in the network being fooled. It was shown then that some deep learning-based face detectors are prone to adversarial attacks not only in a digital domain but also in the real world. In the paper, we investigate the security of the well-known cascade CNN face detection system - MTCNN and introduce an easily reproducible and a robust way to attack it. We propose different face attributes printed on an ordinary white and black printer and attached either to the medical face mask or to the face directly. Our approach is capable of breaking the MTCNN detector in a real-world scenario.

**Index Terms**—adversarial attacks, face detection, MTCNN, physical domain

## I. INTRODUCTION

Contemporary deep learning systems are proved to be almost perfect face detectors, which outperform human abilities in this area [1]. The number of applications in today's life increases tremendously due to this fact. They would replace humans in areas where their accuracy is the most beneficial, for example, security. So since their algorithm-driven decisions may have serious consequences, the question of reliability and robustness against malicious actions becomes crucial. One of this task is face detection which is widely used as preparations operation for FaceID, which allows tracing criminals or control entrance policy.

There are several deep learning approaches to this problem, end-to-end solutions like RetinaNet [1], and cascaded from several NNs like MTCNN [2]. Although end-to-end approach shows better results on synthetic benchmarks, cascaded systems with comparable quality are usually significantly faster.

Unfortunately, there is a technique called adversarial attack, which allows deceiving almost any neural network-based systems in some instances. For example, in the case of a white-box attack in the digital domain, white-box since the

attacker has access to topology and weights on the network, and digital domain because he changes input image pixel-wise. There is no existing solution to mitigate this issue according to recent publications [3] completely. Although these results are interesting from theoretical point of view, in practice, the task of face detection assumes that the processing image is obtained from a real-world device like camera, which is protected from the intrusion, i. e. attacker does not have direct access to the input. It is called physical domain attack. Although there are examples of this type of attacks, they proved to be much harder to perform, since adversarial attacks tend to be very fragile. Insignificant change in environment or illumination usually destroys them. In order to address this issue, a special technique called Expectation-over-Transformation (EoT) was introduced in [4].

In this article, we present the attack on MTCNN face detection system. There are no published attacks on this face detector, although this system is quite well-known and public. Probably, the reason is that this system is robust to adversarial attacks due to its cascaded nature. Since it is hard to use traditional methods (FGSM-like) on the whole system, we decided to attack its first component. It is also worth mentioning that the attacking method implies the use of a public and well-known technique - adversarial attack; the network is available on the Internet and considered to be open, so the work does not violate any law or regulation.

The article is organised as follows. The attack itself is described in section III, the experiments in section IV, and in section II we review related works.

## II. RELATED WORKS

Before describing the proposed method, we review some of the widely used face detection models and their main differences. Then we focus on the adversarial attacks and consider some of the essential works in the area. Since most of the adversarial attacks are applicable in the digital domain and do not pose potential security concerns in applications using

face detection models, we take a closer look at the real-world adversarial attacks and how they can be generated.

### A. Face detection models

The problem of face detection was first practically solved in a seminal work of Viola and Jones [5]. The idea was to apply a hand-crafted Haar feature to the input image pyramid of different scales. Multi-scale pyramid of images allows objects to be detected at different sizes and helps to narrow down the number of proposed regions, thus boosting up the classification. The algorithm is relatively fast, which enables real-time detection but at the cost of poor results with non-frontal faces and low light conditions. The following works [6], [7], [8] mostly focused on improving the proposed architecture.

Recent years have shown that CNN can potentially outperform all classic approaches based on standard features due to its generalization ability. The approach that CNN utilized [9] was quite similar to a classic one: the window with learnt features was sliding over the pyramid of input images and the resulted data fed to a fully-connected layer. Another way of constructing a face detection system was proposed in [10] where authors suggested using the inherent multi-scale, pyramidal structure of DCNN to build feature pyramid.

Apart from the classification mentioned above, face detection networks can be categorized into two classes: single-shot and multi-shot detectors. One of the most well-known examples of the single-shot detector is SSD network [11], which takes an image as an input and computes a feature map with bounding boxes for each class. A similar approach was utilized in [12]. Multi-shot detector (usually two-stage) suggests using several stages. The stages usually include proposal step and refinement steps. One of the most well-known examples of the two-stage detector is R-CNN [13]. The network extracts region proposal with the selected search algorithm and then warps cropped proposals to a square. Features are calculated based on a sparse set of candidates, and the output is fed to a classifier.

The networks with a pure cascade architecture held leading positions for a long time in WIDER FACE challenge [14]. One of the most widely-used was MTCNN detector [2], which performs both face detection and face landmarking. The network uses three sub-networks: P-Net, R-Net and O-Net. The first stage does the coarse face detection producing proposal regions. Then Non-Maximum Suppression algorithm reduces the number of overlapping boxed forming more certain regions, which are fed to R-Net. R-Net refines the selected proposals, and O-Net does the face landmarking. MTCNN is still proposed to be used in the state-of-the-art face recognition system described in [15]. Moreover, it is utilized in the most popular public face recognition implementation of FaceNet [16] available on GitHub<sup>1</sup>.

Besides excellent performance, MTCNN is a promising network in terms of robustness to adversarial attacks. The first

shallow P-Net has a receptive field size of 12x12, which makes the detection resistant to the small artefacts on a face.

### B. Adversarial attacks

Increased interest in neural networks helped to discover not only its state-of-the-art performance in various applications but also its serious drawback concerning robustness against adversarial attacks [17]. Minor and invisible to human eye changes in the input image can mislead the network and make it misclassify the object depicted in the picture with high confidence. It creates a severe security vulnerability in various applications of DCNN.

Suppose a classifier mapping  $f : \mathbb{R}^m \rightarrow \{1\dots k\}$  is given. The loss function is denoted by  $L_f : \mathbb{R}^m \times \{1\dots k\} \rightarrow \mathbb{R}^+$  and is to be continuous. Provided  $x \in \mathbb{R}^m$  is an image and  $l \in \{1\dots k\}$  is a label, the problem can be formulated as follows:

$$\min_{\substack{f(x+r)=l; \\ x+r \in [0,1]^m}} \|r\| \quad (1)$$

In [18] it was proposed to use a box-constrained L-BFGS to find such  $r$ :

$$\min_{x+r \in [0,1]^m} c\|r\| + L_f(x+r, l) \quad (2)$$

As soon as the method for finding adversarial examples was introduced, it was found in [19] that neural networks have a linear nature. This fact enables attacks to be performed in a more computationally efficient way with the use of Fast Gradient Sign Method:

$$X^{adv} = X + \epsilon \text{sign}(\nabla_X J(X, y_{true})) \quad (3)$$

In some cases, one iteration of FGSM might not be sufficient; therefore, FGSM was extended to an iterative version in [20]:

$$X_{N+1}^{adv} = \text{Clip}_{X,\epsilon} \left\{ X_N^{adv} + \alpha \text{sign}(\nabla_X J(X_N^{adv}, y_{true})) \right\} \quad (4)$$

In order to improve the optimization process further on and make the adversarial attacks more robust, momentum was added to iterative FGSM in the paper [21] and the method is usually referred to MI-FGSM. It helps to memorize the gradient direction over iterations, which is a right solution to pass through poor local minima or maxima.

It is commonly assumed, that network weights and model topology are known for an attacker. That is called a white-box attack. It may occur that the model information is not available; thus FGSM cannot be applied; this is called a black-box attack, which is well-explained in [22].

### C. Attacks in physical domain

The adversarial attacks in the physical domain are assumed to be more challenging. The input image fed to the network is subjected to various transforms imposed by the real-world: perspective transform, rotations, and so on. The searching procedure should take into account it and generate such an input that is tolerant of this kind of transforms. Such tolerance is usually achieved with a technique called Expectation

<sup>1</sup><https://github.com/davidsandberg/faceNet>

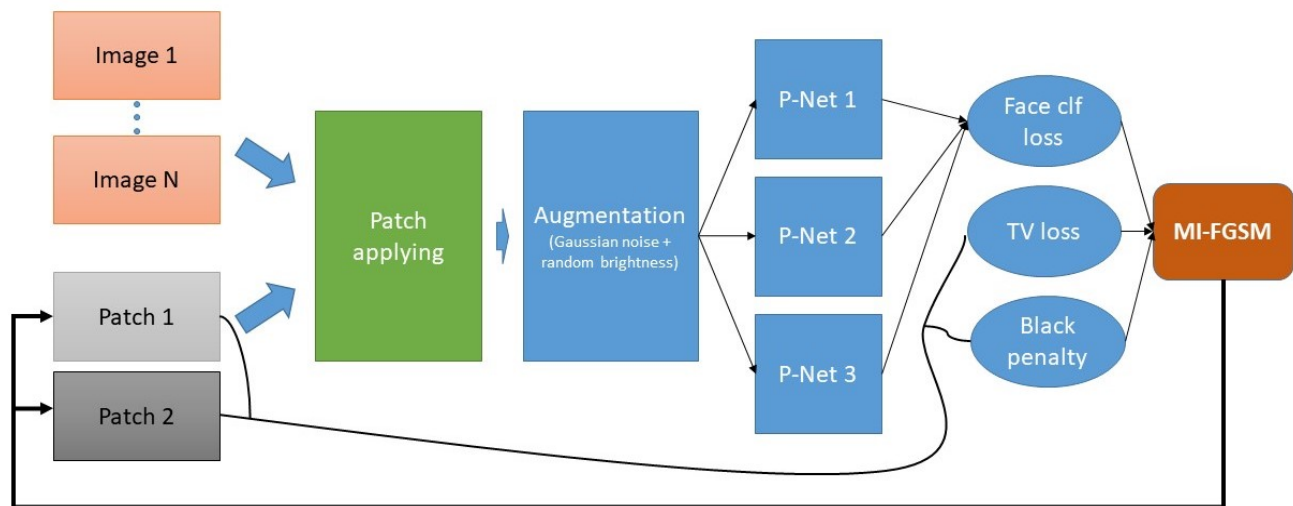


Fig. 1. The attack pipeline. The patches are applied to the batch of images. The resulted images are augmented and fed to three networks. Loss of the face classification output and loss of the patches are computed, and the summed output is back-propagated to the patches.

over Transformation (EoT) [4]. The key point of EoT is to model inherent perturbations in input during the optimization procedure.

Adversarial attacks in the real-world domain might impose more challenging problems as an attacker can mislead the network in a non-intrusive fashion. For instance, with the rise of autonomous vehicles, one can construct an attribute for a stop sign and fool a self-driving car [23]. Authors did another illustrative work in [24], where a particular pattern was generated to avoid detection by a person detector based on YOLO. The patch was trained with various transforms taken into account to enable the real-world attack.

In the article, we focus on a grayscale real face attributes to avoid face detection performed by MTCNN in real-time. The attack is supposed to be a white-box.

### III. PROPOSED METHOD

In this section, we describe the whole process of generating patches. Each sub-network of MTCNN has three output layers: face classification, bounding box regression and face landmark localization. Given this, we came up with three possible approaches for the attack:

- Attacking the face classification layer of P-Net;
- Attacking the bounding boxes layer of P-Net, the similar method for YOLO was described in [25], which exploits a non-secure NMS algorithm;
- Attacking the output layer of O-Net;
- Attacking the whole network.

The first approach compared to the others, requires the least demanding of architecture; hence, the face classification layer P-Net will be used for the attack. Refer to Figure 1 for details of the proposed attack pipeline. In the following subsections, the detailed information on the attack pipeline will be given.

#### A. Expectation-over-Transformation

For adversarial attack it is important to be robust to succeed in the physical domain. This task can be completed via EoT technique which was mentioned before. In our case we performed it in the following way: when an adversarial patch is being trained it does not minimize loss function over a single image, it uses batch consisted of multiple images with different positions of the head instead. Since it minimizes loss over pictures with different size of the patch and different brightness, it should be robust against these types of transforms in the real world. The scheme explaining the process is depicted in Figure 2.

#### B. Projective transform

To apply rectangular patches on different surfaces, we use a projective map. Projective map is defined by its eight coefficients, which can be defined. Firstly in the real world, we label patch location in edges of rectangles. If a patch has curved boundaries, it can be approximated with a rectangular grid. Then coefficients of projective transform are calculated, and the patch is applied. The example of how it is performed is depicted in Figure 3.

#### C. MTCNN analyzer

Once the patches are applied, and the resulted images are augmented, they should be resized to various scales and fed to P-Net. Originally, MTCNN builds up a pyramid of images with a given scale step factor. Using all scales for attacks is not feasible as it is more demanding of a resource. To mitigate this problem, we develop two possible approaches:

- We find the scale contributing most to the detection and use it with up-neighbouring scale and down-neighbouring scale;
- We find the scale that contributes most to the detection and use the scale with a slightly bigger size (which is not

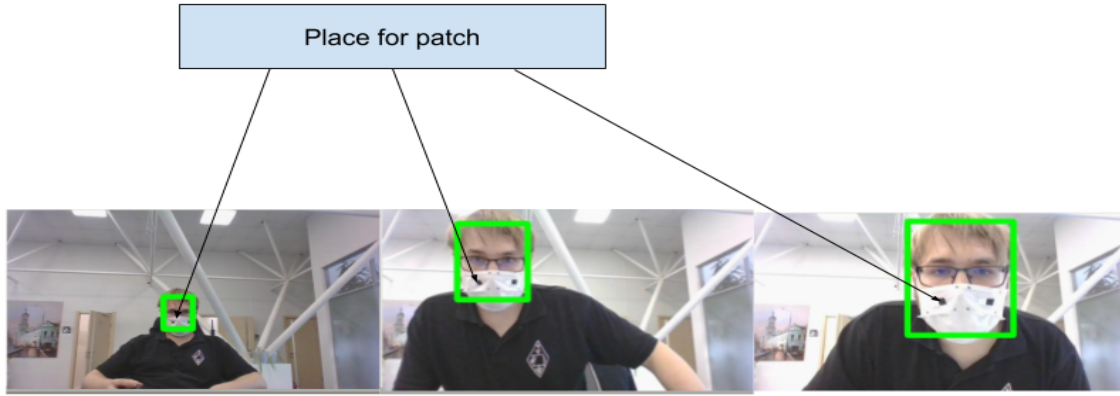


Fig. 2. EoT process. The input images contain different head position and slightly different ambient luminance. The black dots on the mask depict placeholders for the patches.




Fig. 3. Example of a projective map from patches to the face. The face is labelled with printed grid, so for each rectangle coefficients of the transform can be defined.



Fig. 4. The images fed to R-Net. The scales that give the most meaningful pictures for R-Net are 4, 5 and 7. These layers will highly contribute to the refinement stage.

presented in a pyramid originally) and slightly smaller size, i.e. we do a size augmentation.

To find the most contributing scale, we let the image pass through the P-Net and manually trace the scale that gives the most meaningful results to R-Net. The example with pictures size of 24x24 is shown in Figure 4. The more images with face passed to R-Net the more likely the face detection ends up successfully. Once three scales are selected in a way described above, the pyramid is created and loss functions of outputs are calculated.

#### D. Loss functions

The optimization process consisted of two main parts and one optional:

- Face classification loss. The main objective is to lower the probability so that face will not be detected. Two losses were proposed to be used for  $L_{clf}$ :  $L_{inf}$  and  $L_2$ . Both showed comparably good results. As we use three layers, we sum the loss for each scale.
- Total variation loss. To make the optimization give preference for good-looking patterns without sharp color transitions and noise, we calculate  $L_{TV}$  from patches given  $p_{i,j}$  is a pixel value in position  $i, j$ :

$$L_{TV} = \sqrt{(p_{i,j} - p_{i+1,j})^2 + (p_{i,j} - p_{i,j+1})^2} \quad (5)$$

Basically, the smoother transitions the less value of  $L_{TV}$ .

- Black penalty loss. In the case of surgical mask it is reasonable to reduce an amount of black color on the patch, that enables patch to be less unusual. To decrease the black colors, we use  $L_{BLK}$ :

$$L_{BLK} = \sum_{i,j} 1 - p_{i,j} \quad (6)$$

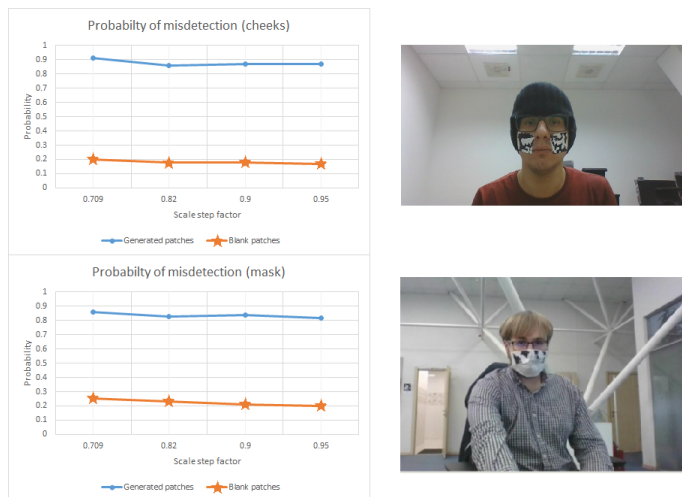


Fig. 5. Experiment results. The probability is averaged over 1000 frames for each scale step factor. With a generated patch, the misdetection probability becomes significantly higher in both cases. The photos on the right are two random frames from the recorded videos.

Finally, to balance the contribution of each optimization coefficients are added and the total loss has the following form:

$$L = \sum_{i=1 \dots 3} L_{clf_i} + \alpha L_{TV} + \beta L_{BLK} \quad (7)$$

All coefficients were derived empirically and should be determined on a case-by-case basis.

#### IV. EXPERIMENTS

To do experiments, we need to define the main parameters for MTCNN. The minimal size parameter (minsize) determining the minimal size of an image in a pyramid was set to be 21 pixels (size of the shortest axis). The thresholds were set to 0.6, 0.7 and 0.7 for each sub-network respectively, and the scale step factor to 0.709. Such parameters are widely used in practice.

The patches were trained for 2000 epochs; further training did not give any improvement. To test a generated pattern, we recorded videos with two setups: a surgical mask with patches and just two patches on cheeks. The videos were used to calculate a misdetection probability for various step scale factors. To make the experiments more valuable in practical sense, the position in videos was different: close-up, mid-range and staying in a far distance. The results are given in Figure 5. It is worth mentioning that the performed attack was targeted, so not well-transferable to other persons not included to the training samples.

#### V. CONCLUSION

In this article we:

- Found the method to attack the popular MTCNN face detector in the digital domain;
- Transferred this attack to the physical domain with EoT technique;

- Verified these results by conducting experiments in the real-world domain.

The obtained results show that one the most robust face detection network is still beyond expectations and need to be improved. The attacks in physical domain impose severe security issues, so possible ways to address it should be found. In the article, we took the first step toward securing the face detection systems by finding a reproducible attack. In the future, we consider attacks using different places of intrusion in MTCNN detection pipeline and think about possible improvement of security.

#### REFERENCES

- [1] J. Deng, J. Guo, Y. Zhou, J. Yu, I. Kotsia, and S. Zafeiriou, "Retinaface: Single-stage dense face localisation in the wild," *ArXiv*, vol. abs/1905.00641, 2019.
- [2] 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, Oct 2016.
- [3] A. Athalye, N. Carlini, and D. Wagner, "Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples," *Proceedings of the 35th International Conference on Machine Learning*, pp. 274–283, 2018.
- [4] A. Athalye, L. Engstrom, A. Ilyas, and K. Kwok, "Synthesizing robust adversarial examples," in *ICML*, 2017.
- [5] P. Viola and M. J. Jones, "Robust real-time face detection," *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, May 2004.
- [6] S. Liao, A. K. Jain, and S. Z. Li, "A fast and accurate unconstrained face detector," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, pp. 211–223, 2016.
- [7] L. Bourdev and J. Brandt, "Robust object detection via soft cascade," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 2, June 2005, pp. 236–243 vol. 2.
- [8] S. C. Brubaker, J. Wu, J. Sun, M. D. Mullin, and J. M. Rehg, "On the design of cascades of boosted ensembles for face detection," *International Journal of Computer Vision*, vol. 77, no. 1, pp. 65–86, May 2008.
- [9] S. Yang, P. Luo, C. C. Loy, and X. Tang, "From facial parts responses to face detection: A deep learning approach," *IEEE International Conference on Computer Vision*, pp. 3676–3684, 2015.
- [10] T.-Y. Lin, P. Dollár, R. B. Girshick, K. He, B. Hariharan, and S. J. Belongie, "Feature pyramid networks for object detection," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 936–944, 2016.
- [11] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, pp. 21–37.
- [12] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," 06 2016, pp. 779–788.
- [13] R. B. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 580–587, 2013.
- [14] S. Yang, P. Luo, C. C. Loy, and X. Tang, "Wider face: A face detection benchmark," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 5525–5533.
- [15] J. Deng, J. Guo, and S. P. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," *ArXiv*, vol. abs/1801.07698, 2018.
- [16] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 815–823, 2015.
- [17] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrđić, P. Laskov, G. Giacinto, and F. Roli, "Evasion attacks against machine learning at test time," in *Proceedings of the 2013th European Conference on Machine Learning and Knowledge Discovery in Databases - Volume Part III*, ser. ECMLPKDD'13. Berlin, Heidelberg: Springer-Verlag, 2013, pp. 387–402.



- [18] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *CoRR*, vol. abs/1312.6199, 2013.
- [19] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *CoRR*, vol. abs/1412.6572, 2014.
- [20] A. Kurakin, I. J. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," *ArXiv*, vol. abs/1607.02533, 2016.
- [21] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, and J. Li, "Boosting adversarial attacks with momentum," *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9185–9193, 2017.
- [22] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, "Practical black-box attacks against machine learning," in *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, ser. ASIA CCS '17. New York, NY, USA: ACM, 2017, pp. 506–519.
- [23] K. Eykholt, I. Evtimov, E. Fernandes, B. Li, A. Rahmati, F. Tramèr, A. Prakash, T. Kohno, and D. Song, "Physical adversarial examples for object detectors," in *Proceedings of the 12th USENIX Conference on Offensive Technologies*, ser. WOOT'18. Berkeley, CA, USA: USENIX Association, 2018, pp. 1–1.
- [24] S. Thys, W. V. Ranst, and T. Goedemé, "Fooling automated surveillance cameras: adversarial patches to attack person detection," *ArXiv*, vol. abs/1904.08653, 2019.
- [25] D. Wang, C. Li, S. Wen, S. Nepal, and Y. Xiang, "Daedalus: Breaking non-maximum suppression in object detection via adversarial examples," *ArXiv*, vol. abs/1902.02067, 2019.