

# Fooling Thermal Infrared Pedestrian Detectors in Real World Using Small Bulbs

Xiaopei Zhu<sup>1</sup>, Xiao Li<sup>2</sup>, Jianmin Li<sup>2</sup>, Zheyao Wang<sup>1\*</sup>, Xiaolin Hu<sup>2\*</sup>

<sup>1</sup>Department of Microelectronics and Nanoelectronics, Tsinghua University, Beijing 100084, China

<sup>2</sup>Department of Computer Science and Technology, Institute for Artificial Intelligence

State Key Laboratory of Intelligent Technology and Systems, BNRist, Tsinghua University, Beijing 100084, China  
{zxp18, lixiao20}@mails.tsinghua.edu.cn, {lijianmin, z.wang, xlhu}@tsinghua.edu.cn

## Abstract

Thermal infrared detection systems play an important role in many areas such as night security, autonomous driving, and body temperature detection. They have the unique advantages of passive imaging, temperature sensitivity and penetration. But the security of these systems themselves has not been fully explored, which poses risks in applying these systems. We propose a physical attack method with small bulbs on a board against the state-of-the-art pedestrian detectors. Our goal is to make infrared pedestrian detectors unable to detect real-world pedestrians. Towards this goal, we first showed that it is possible to use two kinds of patches to attack the infrared pedestrian detector based on YOLOv3. The average precision (AP) dropped by 64.12% in the digital world, while a blank board with the same size caused the AP to drop by 29.69% only. After that, we designed and manufactured a physical board and successfully attacked YOLOv3 in the real world. In recorded videos, the physical board caused AP of the target detector to drop by 34.48%, while a blank board with the same size caused the AP to drop by 14.91% only. With the ensemble attack techniques, the designed physical board had good transferability to unseen detectors.

## Introduction

Deep learning has achieved remarkable success in various tasks such as classification (Karpathy et al. 2014), detection (Redmon and Farhadi 2017), and segmentation (Yin et al. 2020). However, it is well known that deep neural networks (DNNs) are vulnerable to adversarial attacks, i.e., they can be fooled by input examples with some deliberately designed small perturbations. Such examples are called adversarial examples. Since the findings of Szegedy et al. (2013), there is increasing interest in the field of adversarial attacks. For digital world attacks, many methods have been proposed including the gradient-based attacks (Goodfellow, Shlens, and Szegedy 2015; Kurakin, Goodfellow, and Bengio 2017; Madry et al. 2018), optimization-based attacks (Carlini and Wagner 2017; Szegedy et al. 2014; Eykholt et al. 2018), and network-based attacks (Xiao et al. 2018; Liu et al. 2019). Annoyingly, adversarial examples not only exist in the digital world but also exist in the real world. Athalye et al.



Figure 1: An example of physical infrared attack and control experiments

(2018) showed that a 3D printed turtle could be mistaken for a rifle by a DNN. Simen et al. (2019) designed a printable patch that successfully attacked the pedestrian detection system and achieved an “invisibility” effect. Xu et al. (2019) invented Adversarial T-shirts, which could attack the person detectors even if it has non-rigid deformation due to a moving person’s pose changes. The adversarial attack in the physical world has attracted much attention as it poses high risks to the widely deployed deep learning-based security systems. It urges researchers to re-evaluate the safety and reliability of these systems.

Almost all current research on adversarial attacks focused on the visible light field. There is a lack of research on the safety of infrared (Thermal infrared in our paper) object detection systems, which have been widely deployed in our society with their unique advantages. First of all, the infrared object detection systems can work at night. It implies that the surveillance systems based on the infrared cameras do not need environmental light and can save energy in certain scenarios. Some autonomous driving companies are currently using infrared images as auxiliary input at night. Secondly, they can detect objects behind certain obstacles. For example, a person can still be detected when hiding in bushes. Thirdly, compared to visible light images, infrared images not only contain the shape information of the object but also contain the temperature information of the object. During the pandemic caused by COVID-19, infrared pedestrian detection has received more and more attention. With the de-

\*Corresponding Author

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

velopment of deep learning, infrared object detection has made significant progress.

Compared with the visible images which have three channels (RGB), the challenge of the infrared image processing is that the infrared image has only one gray-scale channel, and the texture information is far less than that of visible light image. Besides, to realize physical attacks, visible images can be printed by a laser printer, which can preserve most details of the designed adversarial images. Obviously, one cannot obtain an adversarial infrared image by “printing” any digital image.

To solve this problem, we propose a method to realize the adversarial infrared images in the real world. Our method is to use a set of small bulbs on a cardboard, which can be held by hands. A dedicated circuit is designed to supply electricity to the bulbs. Equipped with eight 4.5V batteries, the cardboard decorated with small bulbs can successfully fool the state-of-the-art infrared pedestrian detection models. The cost of this setup is less than 5 US dollars. An example of physical infrared attack and control experiments is shown in Figure 1. As far as we know, we are the first to realize physical attacks on the thermal infrared pedestrian detector.

## Related Work

### Digital Attacks in Visible Light Images

Szegedy et al. (2014) found that when small perturbations added to an image, it will cause image classification errors. Goodfellow et al. (2015) developed a method to efficiently compute an adversarial perturbation for a given image, which is called FGSM. Based on the above work, BIM (Kurakin, Goodfellow, and Bengio 2017) optimizes the process by taking more small-step iterations. Dong et al. (2018) proposed MIFGSM, which added a momentum term during iterative optimization. A particular case of digital attacks is to modify only one pixel of the image to fool the classifier. Moosavi-Dezfooli et al. (2017) computed universal perturbations to fool DNNs.

### Physical Attacks in Visible Light Images

Sharif et al. (2016) designed a wearable glasses that could attack facial recognition systems. Eykholt et al. (2018) designed a road sign with perturbations that can fool the road sign recognition systems in practical driving. Athalye et al. (2018) proposed a method for generating 3D adversarial examples. Zhou et al. (2018) designed an invisible mask that can attack the face recognition system. They hide the infrared diode under the hat. Simen et al. (2019) proposed an optimization-based method to create a patch that could successfully hide a person from a person detector.

### Object Detection

There are three kinds of target detectors: one-stage, two-stage, and multi-stage detectors. The YOLOv3 (Redmon and Farhadi 2018) model is a typical one-stage detection method with the advantage of fast detection speed and high accuracy. RetinaNet (Lin et al. 2017) is also a one-stage detector using Focal loss. Faster R-CNN (Ren et al. 2017) is a typical two-stage detector using RPN to extract candidate ROI

regions. Cascade R-CNN (Cai and Vasconcelos 2018) is a multi-stage detector which contains a set of detectors trained with increasing IOU thresholds.

Sujoy et al. (2017) used local steering kernel (LSK) as low-level descriptors for detecting pedestrians in thermal infrared images. Yifan et al. (2019) proposed an infrared pedestrian detection method with a converted temperature map. YuHang et al. (2020) proposed an infrared small target detection algorithm based on the peak aggregation number and Gaussian discrimination. Mate et al. (2020) investigated the task of automatic person detection in thermal images using convolutional neural network models originally intended for detection in RGB images. They compared the performance of the standard state-of-the-art object detectors such as Faster R-CNN, SSD, Cascade R-CNN, and YOLOv3, that were retrained on a dataset of thermal images extracted from videos. Videos were recorded at night in clear weather, rain, and in the fog, at different ranges, and with different movement types. YOLOv3 was significantly faster than other detectors while achieving performance comparable with the best (Kristo, Ivacic-Kos, and Pobar 2020). In this paper, we mainly attacked the YOLOv3 model. After that, we transferred the infrared patch attack to other object detectors.

## Generating Infrared Adversarial Patches

We first formulate the problem, and introduce the attack method in the digital world, and then the attack method in the physical world.

### Problem Formulation

We assume that the original image input is  $x$ , and the adversarial perturbation is  $\delta$ . Since we use a patch attack method, the perturbation only occupies part of the image. We assume the patched image is  $\tilde{x}$ .

Let  $f$  denote a model,  $\theta$  denote its parameters, and  $f(x, \theta)$  denote its output given the input  $x$ . Note that most object detector models have three outputs: position of the bounding box  $f_{pos}(x, \theta)$ , the object probability  $f_{obj}(x, \theta)$ , and the class score  $f_{cls}(x, \theta)$ . Our goal is to attack the detection model so that the detection model can not detect objects of the specified category. In other words, we want to lower the  $f_{obj}(x, \theta)$  score as much as possible. Therefore, the goal can be described as

$$\min_{\delta} f_{obj}(\tilde{x}, \theta). \quad (1)$$

Our goal is to attack the object detector in the physical world, so we need to consider various image transformations of the patches during the attack, such as rotation, scale, noise, brightness, and contrast. Furthermore, due to lots of intra-class variety of pedestrians, we hope to achieve a universal attack on different people. Assuming that the set of transformations is  $T$ , the patched image considering patch transformations is  $\tilde{x}_t$ , and the data set has  $N$  pedestrians, the goal can be described as

$$\min_{\delta} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{t \in T} f_{obj}^{(i)}(\tilde{x}_t, \theta). \quad (2)$$

Our loss function consists of two parts:

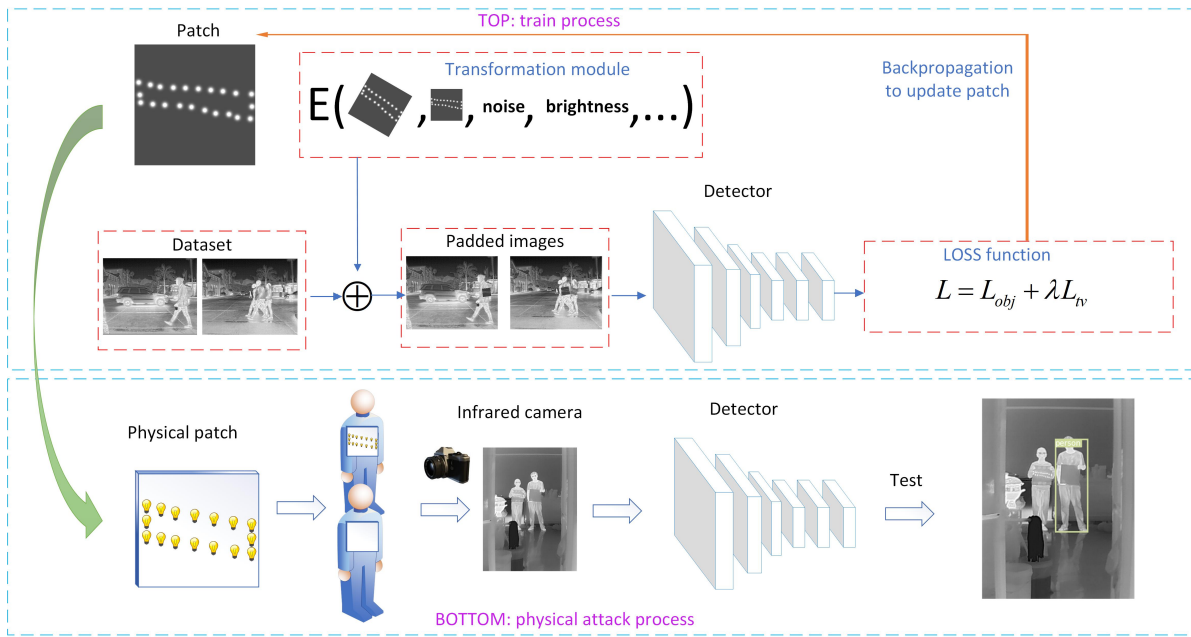


Figure 2: Overview of our work. Top is the training process, bottom is the physical attack process.

- $L_{obj}$  represents the maximum objectness score as shown in Equation (2).
- $L_{tv}$  represents the total variation of the image. This loss ensures that the optimized image is smoother and prevents noisy images. We assume that  $p_{i,j}$  represents the pixel value at a coordinate  $(i, j)$  in the patch. We can calculate  $L_{tv}$  as follows:

$$L_{tv} = \sum_{i,j} \sqrt{(p_{i,j} - p_{i+1,j})^2 + (p_{i,j} - p_{i,j+1})^2}. \quad (3)$$

We take the sum of the two losses weighted by the factor  $\lambda$  which are determined empirically. Given by:

$$L = L_{obj} + \lambda L_{tv}. \quad (4)$$

For the ensemble attack, We hope to lower the maximum objectness score of each detector at the same time. Assume there are  $M$  detectors, and the maximum objectness score of  $i$ -th detector is  $L_{obj}^{(i)}$ . We take the sum of these losses. So the total loss of the ensemble attack is:

$$L_{ensemble} = \sum_{i=1}^M L_{obj}^{(i)} + \lambda L_{tv}. \quad (5)$$

We use backpropagation to update the the patch iteratively (Figure 2).


## Digital Patch Attack

**The Pixel-Level Patch** First of all, we wondered if we follow the adversarial attack method on visible light images to fool pedestrian detectors, what patch will be resulted in. Specifically, can the resulted patches be realized easily by

using some thermal materials? It is easy to carry out this experiment because we only need to change the RGB images to grayscale images and follow the method described in (Thys, Ranst, and Goedemé 2019).

**The Gaussian Functions Patch** To implement the adversarial attack in the physical world, another idea is to design an adversarial example based on the thermal property of certain given electronic components (e.g., resistors). We can first measure the relationship between the thermal properties of the components and the image patterns captured by the infrared cameras, then design an adversarial patch in the digital world, and finally manufacture a physical board specified by this digital adversarial patch. Since infrared thermal imaging mainly uses the thermal radiation of the object, in the selection of electronic components, we consider diodes, resistors and small bulbs. We found that the small bulb is a good candidate for adjusting image patterns captured by infrared cameras. Its brightness well reflects its temperature. With the help of the rheostat, we can adjust the bulb brightness intuitively. We can fine-tune its infrared imaging pattern in this way.

We took an infrared image of a single bulb. Then we used FLIR tools software provided by FLIR company to export the temperature of each point in the image. We first selected the temperature values on multiple lines that cross the center and then we used the Gaussian function for fitting as shown in Figure 3. The fitting was good and the Root Mean Squard Error (RMSE) was 0.1511. The amplitude value of the Gaussian function is 10.62, and the standard deviation is 70.07. Further experiments showed that the temperature of the same point measured by the infrared camera did not change with distance because of the correction function of the infrared camera. Therefore, the pixel value of the same



(a) Image of a single bulb (b) Temperature fitting

Figure 3: The measured temperatures along the horizontal section shown in (a) are plotted in (b). The temperatures along other sections were nearly the same as those along the horizontal section, therefore are not shown in (b).

point did not change with distance.

If we put many bulbs on a cardboard, the infrared camera will capture an image patch with a set of 2D Gaussian patterns, the centers of which have the highest pixel values. The problem now is whether we can design such an image patch to fool the pedestrian detectors. We first carry out the attack using the Gaussian functions patch in the digital world. Since the amplitude and standard deviation of the Gaussian function is fixed to be measured values, the optimization parameter of each two-dimensional Gaussian function is the coordinate of the center point. Besides, this kind of patch significantly reduces the number of parameters compared with the pixel-level patch. In our experiment, the number of optimization parameters dropped nearly 1000 times.


Assuming that the pattern of a patch is superimposed by  $M$  spots that conform to Gaussian functions, where the center point of the  $i$ -th Gaussian function is  $(p_x, p_y)$ , the amplitude amplification factor is  $s_i$ , and the standard deviation is  $\sigma_i$ . The measured  $s_i$  was 10.62, and  $\sigma_i$  was 70.07 in our experiment. We assume that the height of the entire image is  $h$ , the width is  $w$ , and the coordinate of a single-pixel is  $(x, y)$ , where  $x \in [0, w], y \in [0, h]$ , then the  $i$ -th Gaussian function is as follows:

$$g^{(i)}(x, y) = s_i \cdot \exp\left(-\frac{(x - p_x^{(i)})^2 + (y - p_y^{(i)})^2}{2\sigma_i^2}\right). \quad (6)$$

Suppose the background of the patch is  $P_{back}$ , which is a matrix with all elements equal to  $\mu$ . The overall function of multiple 2D Gaussian functions superimposed together is denoted by  $P_{syn}$ :

$$P_{syn} = P_{back} + \sum_{i=1}^M G_i \quad (7)$$

$$G_i = \begin{pmatrix} g^{(i)}(0, 0) & \dots & g^{(i)}(0, w) \\ \vdots & \ddots & \vdots \\ g^{(i)}(h, 0) & \dots & g^{(i)}(h, w) \end{pmatrix}. \quad (8)$$



(a) Movable position arrangement (b) Circuit design (demo)

Figure 4: Physical board design

## Physical Board Attack

In practice, we face a challenge to move the bulbs freely on a board when we try different patterns. In other words, we need an adjustable physical board as shown in Figure 4(a). We solve the problem by using magnets. One magnet is fixed to the bulb, and the other magnet is placed on the other side of the board. The magnet can attract the bulb on the other side and can adjust the position of the bulb as a button.

For circuit design, we used multiple independent power supply DC4.5V power modules. The rated voltage of the small bulb was 3.8V. After measurement, the total power of the physical board did not exceed 22W. The power supply lines of each light bulb were connected in parallel, and each circuit contained a small switch and a rheostat. This can ensure that the different small bulbs were independent of each other. The demo circuit design diagram is shown in Figure 4(b).

## Experiments and Results

### Preparing the Data


The dataset we used is FLIR\_ADAS dataset v1.3 released by FLIR. FLIR\_ADAS provides an annotated thermal image and non-annotated RGB image set for training and validation of object detection networks. The data set contains 10228 thermal images sampled from short videos and a continuous 144-second video. These photos were taken on the streets and highways in Santa Barbara, California, the USA from November to May. The thermal image is a FLIR Tau2 (13 mm f/1.0, 45-degree HFOV and 37-degree VFOV, FPA640 × 512, NETD < 60mK).

Thermal images are manually annotated and contain a total of four types of objects, namely *people*, *bicycles*, *cars*, and *dogs*. Since we only care about *people*, we filtered the original data set, and only kept images that contained people whose height is greater than 120 pixels. We finally selected 1011 images containing infrared pedestrians. We used 710 of them as the training set and 301 as the test set. We named them *FLIR\_person.select*.

### Target Detector

Mate et al.(2020) have compared the performance of the standard state-of-the-art infrared object detectors such as





(a) The pixel-level patch (b) The Gaussian functions patch

Figure 5: The digital patch

Faster R-CNN, SSD, Cascade R-CNN, and YOLOv3. They found that YOLOv3 was significantly faster than other detectors while achieving performance comparable with the best. So we chose YOLOv3 as the target detector. The network has 105 layers. We resized the input images to  $416 \times 416$  as required by the model. We chose the pre-training weights officially provided by YOLO and then finetuned on *FLIR\_person\_select*. The AP of the model was 0.9902 on the training set, and 0.8522 on the test set. We used this model as the target model of attack.

### Simulation of Physical Attacks

**Pixel-Level Patch Attack** Following the process described by Simen et al. (2019), we obtained a patch shown in Figure 5 (a). The attack was successful as the patch made the accuracy of YOLOv3 dropped by 74.57% (see Supplementary Information for more details). However, the resulted patch contained numerous grayscale pixels as noise, which are difficult to realize physically. Therefore we abandoned this approach.

**Gaussian Functions Patch Attack** The pattern of the Gaussian functions patch is superimposed by multiple spots that conform to a two-dimensional Gaussian function (Figure 3). To make the patch more robust, we designed a variety of transformations including random noise on the patch, random rotation of the patch (clockwise or counterclockwise within 20 degrees), random translation of the patch, and ran-




Figure 6: Evaluation of Gaussian functions patch attack

The number	The AP dropped by
9	46.02%
15	51.26%
22	64.12%
25	65.74%
36	66.88%

Table 1. The study of the number of Gaussian functions

dom changes in the brightness and contrast of the patch. These transformations simulate the perturbation of the physical world to a certain extent, which effectively improves the robustness of the patch. Then we used the training set of *FLIR\_person\_select* and placed the patch on the upper body of the pedestrians according to the position of the bounding box. The size of the patch was 1/5 of the height of the bounding box. Next, we used the patched image as input and ran the YOLOv3 pedestrian detector we had trained. We used a stochastic gradient descent optimizer with momentum, and the size of each batch was 8. The optimizer used the back-propagation algorithm to update the parameters of Gaussian functions by minimizing Equation (4). Through this process, we obtained a series of patches with different numbers of Gaussian functions. Figure 5(b) is an example with 22 Gaussian functions.

Next, we applied the optimized patch which is shown in Figure 5(b) to the test set, using the same process we used during training, including various transformations. We used random noise patches with maximum amplitude value 1 and constant pixel value patches (blank patches) for control experiments. The pixel values of blank patch in our experiment were 0.75. We tried other values and found that blank patches with different pixel values had a similar attack effect. We applied these different patches to the *FLIR\_person\_select* test set, and then input the patched images to the same detection network to test its detection performance. We adopt the IOU method to calculate the accuracy of the detection. The precision-recall (PR) curves are shown in Figure 6. Using the output of the clean image input as ground truth, the Gaussian functions patch we designed made the average precision (AP, the area under the PR curve) of the target detector drop by 64.12%. We give an example to show the attacking effect in Figure 7. In contrast, the AP of the target detector dropped by 25.05% and 29.69% using random noise patch and using blank patch, respectively.

Note that the attack performance of the Gaussian functions patch attack was not as good as the pixel-level patch attack. This is reasonable as the latter had nearly 1000 times more parameters than the former. But the former is easier to be realized physically.

We tried different kinds of patches and evaluated the attack effect of different patches by AP, the results is shown in Table 1. We found the Gaussian functions patch with 22 spots had a good attack effect, while maintaining a small amount of parameters.

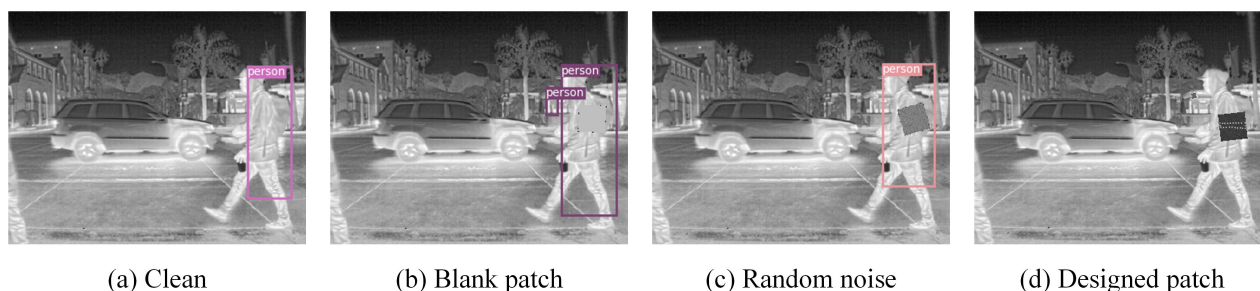


Figure 7: Gaussian functions patch attack and control experiments

### The Effects of Patch Size

We scaled up or down the original image to study the effect of patch size on the attack. We used the patch shown in Figure 5(b). We did five experiments. One kept the original size of the patch ( $300 \times 300$ ), the other two expanded the side length by 1.5 times and 2 times respectively, and the last two reduced the side length to  $2/3$  and  $1/2$  of the original respectively. The results are shown in Figure 8. The patch which is doubled in size caused AP of YOLOv3 to drop by 95.42%. We found that when the patch size dropped to  $1/2$  of its original size, its attack performance dropped a lot. That's the limit of our patch attack method.

### Evaluation of Attacks in the Real World

The pattern of the physical board is derived from the Gaussian functions patch as shown in Figure 5(b). We chose a  $35\text{cm} \times 35\text{cm}$  cardboard. The finished board is shown in Figure 9. It is worth noting that the total manufacturing cost of our physical board did not exceed \$5, indicating the proposed approach is economic. Figure 10 shows the simulated and actual boards.

We conducted physical attack experiments. The equipment we used was HTI-301 infrared camera (FPA  $384 \times 288$ , NETD  $< 60\text{mK}$ ). We invited several people to do the experiment. They could hold the adversarial board, or a blank board, or nothing. We used the infrared camera to shoot these people at the same time under the same conditions, and then sent the thermal infrared photos to the pedestrian detector for testing. Some examples of testing results are shown in Figure 11. It is seen that whenever a person held the blank

board or nothing, YOLOv3 could detect her/him. However, if a person held the adversarial board, YOLOv3 could not detect her/him. The results show the physical board we designed can successfully attack the infrared pedestrian detection system in the real world.

To quantify the effects of a physical attack, we recorded 20 videos from different scenes (See supplementary materials for the demo video). We invited several people to be the actors of the videos. For fair comparison, we asked the actors to walk three times from the same starting position to the same end positions with the same path, once with the adversarial board, once with the blank board and another with nothing. Each group of videos were taken in the same condition. Ten videos were recorded indoors, and the others were recorded outdoors. Each video took 5-10 seconds. The camera got 30 frames per second. We considered different distances (between 3 and 15 meters) and angles (from




Figure 8: Evaluation of patch attack with different size





Figure 9: Physical board we manufactured



(a) The simulation patch (b) The actual patch

Figure 10: Comparison between simulated and actual board



Figure 11: Physical board attacks and control experiments

45 degrees on the left to 45 degrees on the right) when recording the video. There were 1853 frames of images that contained the physical board we designed, 1853 frames of images that contained the blank board, and the other 1853 frames that just contained pedestrians. The result showed that the cardboard caused the AP of the target detector to drop by 34.48%, while a blank board with the same size caused the AP to drop by 14.91% only.


### Ensemble Attack

In order to study whether we can transfer the infrared patch attack to other detection models, we did the following experiment. At the beginning, we directly used the patch which successfully attacked YOLOv3 to attack other detectors, such as Cascade-RCNN and RetinaNet. The patch trained on YOLOv3 caused the AP of Cascade-RCNN and RetinaNet to drop by 11.60% and 25.86%. To improve the transferability of attack, we use model ensemble techniques. We obtained a new Gaussian patch by integrating YOLOv3, Faster-RCNN, and Mask-RCNN during training as shown in Figure 12. Cascade-RCNN and RetinaNet were attacked in the

digital world firstly. As shown in Table 2, our patch caused the AP of Cascade RCNN and RetinaNet to drop by 35.28% and 46.95% respectively. After that, we conducted experiments in the physical world, and the experimental settings were consistent with section *Evaluation of attacks in the real world*. Our patch make the AP of Cascade RCNN and RetinaNet dropped by 21.67% and 25.17% respectively in real world.

### Conclusion

In this article, we demonstrate that it is possible to attack the infrared pedestrian detector in the real world. We propose two kinds of patches: the pixel-level patch and the Gaussian functions patch. We implement the attack in the physical world by designing a cardboard decorated with small bulbs. The physical board can successfully fooled the infrared pedestrian detector based on YOLOv3. In addition, by using the ensemble attack technique, we designed a cardboard that could fool detectors that were unknown to us. As the thermal infrared detection systems are widely used in night security, automatic driving, especially body temperature detection during COVID-19, our work has important practical significance.



(a) Digital patch of ensemble attack (b) Physical board of ensemble attack

Figure 12: The patch obtained by model ensemble

Train \ Test	Cascade RCNN	RetinaNet
	YOLOv3	11.60%
YOLOv3+Faster-RCNN+Mask-RCNN	35.28%	46.95%

Table 2. Transferability in the digital world

## Acknowledgements

This work was supported in part by National Natural Science Foundation of China under Grant 61734004, Grant U19B2034, Grant 62061136001 and Grant 61620106010.

## References

- Athalye, A.; Engstrom, L.; Ilyas, A.; and Kwok, K. 2018. Synthesizing Robust Adversarial Examples. In Dy, J. G.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, 284–293. PMLR. URL <http://proceedings.mlr.press/v80/athalye18b.html>.
- Biswas, S. K.; and Milanfar, P. 2017. Linear Support Tensor Machine With LSK Channels: Pedestrian Detection in Thermal Infrared Images. *IEEE Trans. Image Process.* 26(9): 4229–4242. doi:10.1109/TIP.2017.2705426. URL <https://doi.org/10.1109/TIP.2017.2705426>.
- Cai, Z.; and Vasconcelos, N. 2018. Cascade R-CNN: Delving Into High Quality Object Detection. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, 6154–6162. IEEE Computer Society. doi:10.1109/CVPR.2018.00644. URL [https://openaccess.thecvf.com/content\\_cvpr\\_2018/papers/Cai\\_Cascade\\_R-CNN\\_Delving\\_CVPR\\_2018\\_paper.pdf](https://openaccess.thecvf.com/content_cvpr_2018/papers/Cai_Cascade_R-CNN_Delving_CVPR_2018_paper.pdf).
- Carlini, N.; and Wagner, D. A. 2017. Towards Evaluating the Robustness of Neural Networks. In *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*, 39–57. IEEE Computer Society. doi:10.1109/SP.2017.49. URL <https://doi.org/10.1109/SP.2017.49>.
- Dong, Y.; Liao, F.; Pang, T.; Su, H.; Zhu, J.; Hu, X.; and Li, J. 2018. Boosting Adversarial Attacks With Momentum. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, 9185–9193. IEEE Computer Society. doi:10.1109/CVPR.2018.00957. URL <http://arxiv.org/abs/1710.06081>.
- Eykholt, K.; Evtimov, I.; Fernandes, E.; Li, B.; Rahmati, A.; Xiao, C.; Prakash, A.; Kohno, T.; and Song, D. 2018. Robust Physical-World Attacks on Deep Learning Visual Classification. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, 1625–1634. IEEE Computer Society. doi:10.1109/CVPR.2018.00175. URL <http://arxiv.org/abs/1707.08945>.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2015. Explaining and Harnessing Adversarial Examples. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. URL <http://arxiv.org/abs/1412.6572>.
- Jiang, Y.; Dong, L.; Chen, Y.; and Xu, W. 2020. An Infrared Small Target Detection Algorithm Based on Peak Aggregation and Gaussian Discrimination. *IEEE Access* 8: 106214–106225. doi:10.1109/ACCESS.2020.3000227. URL <https://doi.org/10.1109/ACCESS.2020.3000227>.
- Karpathy, A.; Toderici, G.; Shetty, S.; Leung, T.; Sukthankar, R.; and Li, F. 2014. Large-Scale Video Classification with Convolutional Neural Networks. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, 1725–1732. IEEE Computer Society. doi:10.1109/CVPR.2014.223. URL <https://doi.org/10.1109/CVPR.2014.223>.
- Kristo, M.; Ivacic-Kos, M.; and Pobar, M. 2020. Thermal Object Detection in Difficult Weather Conditions Using YOLO. *IEEE Access* 8: 125459–125476. doi:10.1109/ACCESS.2020.3007481. URL <https://doi.org/10.1109/ACCESS.2020.3007481>.
- Kurakin, A.; Goodfellow, I. J.; and Bengio, S. 2017. Adversarial Machine Learning at Scale. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. URL <https://openreview.net/forum?id=BJm4T4Kgx>.
- Lin, T.; Goyal, P.; Girshick, R. B.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, 2999–3007. IEEE Computer Society. doi:10.1109/ICCV.2017.324. URL <https://doi.org/10.1109/ICCV.2017.324>.
- Liu, A.; Liu, X.; Fan, J.; Ma, Y.; Zhang, A.; Xie, H.; and Tao, D. 2019. Perceptual-Sensitive GAN for Generating Adversarial Patches. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, 1028–1035. AAAI Press. doi:10.1609/aaai.v33i01.33011028. URL <https://doi.org/10.1609/aaai.v33i01.33011028>.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. URL <https://openreview.net/forum?id=rJzIBfZAb>.
- Moosavi-Dezfooli, S.; Fawzi, A.; Fawzi, O.; and Frossard, P. 2017. Universal Adversarial Perturbations. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, 86–94. IEEE Computer Society. doi:10.1109/CVPR.2017.17. URL <https://doi.org/10.1109/CVPR.2017.17>.
- Redmon, J.; and Farhadi, A. 2017. YOLO9000: Better, Faster, Stronger. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, 6517–6525. IEEE Computer Society. doi:10.1109/CVPR.2017.690. URL <https://doi.org/10.1109/CVPR.2017.690>.
- Redmon, J.; and Farhadi, A. 2018. YOLOv3: An Incremental Improvement. *CoRR* abs/1804.02767. URL <http://arxiv.org/abs/1804.02767>.



Ren, S.; He, K.; Girshick, R. B.; and Sun, J. 2017. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 39(6): 1137–1149. doi:10.1109/TPAMI.2016.2577031. URL <https://doi.org/10.1109/TPAMI.2016.2577031>.

Sharif, M.; Bhagavatula, S.; Bauer, L.; and Reiter, M. K. 2016. Accessorize to a Crime: Real and Stealthy Attacks on State-of-the-Art Face Recognition. In Weippl, E. R.; Katzenbeisser, S.; Kruegel, C.; Myers, A. C.; and Halevi, S., eds., *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, 1528–1540. ACM. doi:10.1145/2976749.2978392. URL <https://doi.org/10.1145/2976749.2978392>.

Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I. J.; and Fergus, R. 2014. Intriguing properties of neural networks. In Bengio, Y.; and LeCun, Y., eds., *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*. URL <http://arxiv.org/abs/1312.6199>.

Thys, S.; Ranst, W. V.; and Goedemé, T. 2019. Fooling Automated Surveillance Cameras: Adversarial Patches to Attack Person Detection. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2019, Long Beach, CA, USA, June 16-20, 2019*, 49–55. Computer Vision Foundation / IEEE. doi:10.1109/CVPRW.2019.00012. URL <http://arxiv.org/abs/1904.08653>.

Xiao, C.; Li, B.; Zhu, J.; He, W.; Liu, M.; and Song, D. 2018. Generating Adversarial Examples with Adversarial Networks. In Lang, J., ed., *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, 3905–3911. [ijcai.org](http://ijcai.org). doi:10.24963/ijcai.2018/543. URL <https://doi.org/10.24963/ijcai.2018/543>.

Xu, K.; Zhang, G.; Liu, S.; Fan, Q.; Sun, M.; Chen, H.; Chen, P.; Wang, Y.; and Lin, X. 2019. Evading Real-Time Person Detectors by Adversarial T-shirt. *CoRR* abs/1910.11099. URL <http://arxiv.org/abs/1910.11099>.

Yin, Z.; Yiu, V.; Hu, X.; and Tang, L. 2020. End-to-End Face Parsing via Interlinked Convolutional Neural Networks. *Cognitive Neurodynamics* doi:<https://doi.org/10.1007/s11571-020-09615-4>.

Zhao, Y.; Cheng, J.; Zhou, W.; Zhang, C.; and Pan, X. 2019. Infrared Pedestrian Detection with Converted Temperature Map. In *2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference, APSIPA ASC 2019, Lanzhou, China, November 18-21, 2019*, 2025–2031. IEEE. doi:10.1109/APSIPAASC47483.2019.9023228. URL <https://doi.org/10.1109/APSIPAASC47483.2019.9023228>.

Zhou, Z.; Tang, D.; Wang, X.; Han, W.; Liu, X.; and Zhang, K. 2018. Invisible Mask: Practical Attacks on Face Recognition with Infrared. *CoRR* abs/1803.04683. URL <http://arxiv.org/abs/1803.04683>.