

# Physically Adversarial Infrared Patches with Learnable Shapes and Locations

Xingxing Wei\*, Jie Yu, Yao Huang

Institute of Artificial Intelligence, Beihang University, Beijing, 100191, China

xxwei@buaa.edu.cn, sy2106137@buaa.edu.cn, y\_huang@buaa.edu.cn

## Abstract

Owing to the extensive application of infrared object detectors in the safety-critical tasks, it is necessary to evaluate their robustness against adversarial examples in the real world. However, current few physical infrared attacks are complicated to implement in practical application because of their complex transformation from digital world to physical world. To address this issue, in this paper, we propose a physically feasible infrared attack method called "adversarial infrared patches". Considering the imaging mechanism of infrared cameras by capturing objects' thermal radiation, adversarial infrared patches conduct attacks by attaching a patch of thermal insulation materials on the target object to manipulate its thermal distribution. To enhance adversarial attacks, we present a novel aggregation regularization to guide the simultaneous learning for the patch's shape and location on the target object. Thus, a simple gradient-based optimization can be adapted to solve for them. We verify adversarial infrared patches in different object detection tasks with various object detectors. Experimental results show that our method achieves more than 90% Attack Success Rate (ASR) versus the pedestrian detector and vehicle detector in the physical environment, where the objects are captured in different angles, distances, postures, and scenes. More importantly, adversarial infrared patch is easy to implement, and it only needs 0.5 hours to be constructed in the physical world, which verifies its effectiveness and efficiency.

## 1. Introduction

Deep Neural Networks (DNNs) have shown promising performance in various vision tasks, including object detection [14], classification [8], face recognition [16], and autonomous driving [15]. However, it is typically known that DNNs are vulnerable to adversarial examples [2, 5, 26], i.e., the human-imperceptible perturbed inputs can fool the DNNs-based system to give wrong predictions. Moreover,




Figure 1. The generation process of adversarial infrared patches. We see the pedestrian cannot be detected after the infrared patches are pasted on the pedestrian in the physical world.

these adversarial examples can be exploited in the physical world. In such cases, a widely used technique is called adversarial patches [1, 4, 24, 28], which have been successfully applied to traffic sign detection by generating a carefully designed sticker [4, 23], or face recognition by adding specific textures on eyeglass frames [17, 22]. The success of adversarial patches raises the concerns because of their great threat to the deployed DNN-based systems in the real world.

Nowadays, famous for its strong anti-interference ability in the severe environment, object detection in the thermal infrared images has been widely used in many safety-critical tasks such as security surveillance [19], remote sensing [25], etc. Consequently, it is necessary to evaluate the physical adversarial robustness of infrared object detectors. However, the aforementioned adversarial patches cannot work well in the infrared images because they depend on the adversarial perturbations generated from the view of RGB appearance. These perturbations cannot be captured by infrared cameras, which perform the imaging by encoding the objects' thermal radiation [20]. Although few recent works have been proposed to address this issue, they have their own limitations. For example, adversarial bulbs based on a cardboard of alight small bulbs [30] are complicated to implement in the real world, and are also not stealthy because they produce heat source. Adversarial clothing [29] based on a large-scale QR code improves the stealthiness by utilizing the thermal insulation material to cover the object's thermal radiation, but still has complex transformations from digital to physical world, making it not easy to implement in the real world.

\*Corresponding author

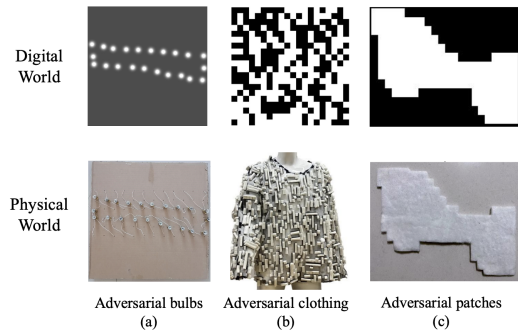


Figure 2. The comparison between different infrared attacks.

In this paper, we propose a physically stealthy and easy-to-implement infrared attack method called “adversarial infrared patches”. Considering the imaging mechanism of infrared cameras by capturing objects’ thermal radiation, we also attach the thermal insulation materials on the target object to manipulate its thermal distribution to ensure the stealthiness. But different from adversarial clothing [29] via the complex QR code pattern, we utilize a simple patch to crop the thermal insulation material, and then adjust the patch’s shape and location on the target object to conduct attacks. Compared with adversarial RGB perturbations, the changes of shapes and locations of the thermal patch can be accurately captured by infrared cameras, which helps perform an effective attack.

However, the shape and location are two kinds of different variables, it is challenging to directly optimize them via unified gradients. For that, we present a novel aggregation regularization to guide the simultaneous learning for the patch’ shape and location on the target object. Specifically, an aggregation degree is proposed to quantify how close one pixel’s neighbours are to being a clique. By combining this metric with the attack goal, the object’s pixels needing to be covered by the thermal insulation material will automatically be gathered to form a valid shape on the available location of the object. In this way, we can adapt a simple gradient-based optimization to solve for the optimal shape and location of the patch. An example of our adversarial infrared patch against the pedestrian detector is shown in Figure 1, and a comparison with the existing physical infrared attacks is given in Figure 2. We see that adversarial infrared patch is simpler than other methods in the digital world, and we just need to crop the thermal insulation materials according to the learned shape, and then paste the patch on the learned location of the pedestrian, which is more easy-to-implement than other methods in the real world.

Our contributions can be summarized as follows:

- We propose the novel “adversarial infrared patches”, a physically stealthy and easy-to-implement attack method for infrared object detection. Instead of generating adversarial perturbations, we perform attacks by learning the available shape and location on the tar-

get object. Owing to this careful design, adversarial infrared patches are easier to implement in the physical world than the existing methods.

- We design a novel aggregation regularization to guide the simultaneous learning for the patch’ shape and location on the target object. Thus, a simple gradient-based optimization can be adapted to solve for the optimal shape and location of the patch.
- We verify the adversarial infrared patches in the pedestrian detection task from both the digital and physical world. Experiments show that adversarial infrared patches can work well in various angles, distances, postures, and scenes, and achieve competitive attacking performance with the SOTA infrared attack while only costing five percent of their time to construct physical adversarial examples. We also extend our method to the vehicle detection task to verify its generalization.

## 2. Related Works

Adversarial examples exist widely in various fields. Recently, adversarial examples in infrared images are explored. Edwards, *et al.* [3] investigate the performance of adversarial attack in ship detection under thermal infrared images. [9] explores how to generate visually imperceptible adversarial infrared examples where an object detector based on DNN cannot detect the targets in the image. These methods generate perturbation by changing pixel values within the infrared image, and thus are not applicable in the physical world.

To address this issue. Zhu, *et al.* [30] make the first attempt to generate a physical adversarial example using a set of small bulbs. These small bulbs change the infrared radiation distribution of the object via simulating the extra heat sources. Although the decorated bulbs can fool infrared pedestrian detectors, they are not stealthy and easy-to-implement. The reason for poor stealthiness is that they produce heat source rather than cover them, which is not suitable for some scenes that need to hide heat sources from the thermal detector. The reason for poor implementation is that they need to carefully design a circuit board to light the small bulbs, which is complicated in real world. Afterwards, Zhu, *et al.* [29] propose adversarial clothing which wraps the whole body to fool the infrared detector at different angles. They aim to design the QR code pattern in the clothing and then transfer this 2D clothing to 3D clothing. The main issue is the complex operations from digital world to physical world. Specifically, they first print the QR code on a large-scale cloth. Then, a tailor is hired to make the cloth into a piece of clothing. Finally, the infrared material is cropped into blocks and pasted on the black area of the clothes.

Our method is different from the existing methods as follows: (1) Compared with [29], our adversarial infrared patches just need to design the shape and location of infrared

patches rather than the complex patterns. Therefore, it is easy and efficient to implement in the physical world. (2) Because the thermal insulation materials can reduce the pedestrian’s thermal radiation rather than enhance it like [30], our method is stealthy in the practical scene. (3) [29] and [30] only verify their methods in the pedestrian detection task, while we verify the proposed methods in both the pedestrian detection and vehicle detection, showing the good generalization.

### 3. Methodology

For simplicity, we introduce our method in the pedestrian detection task. The extension to other task is straightforward.

#### 3.1. Problem Formulation

In the infrared pedestrian detection, suppose  $f(\cdot)$  denotes the pedestrian detector with parameter  $\theta$ , and  $f(\mathbf{x}; \theta)$  denotes the output prediction when a clean thermal image  $\mathbf{x} \in \mathbb{R}^{h \times w}$  is given. For simplicity, we here assume there is only one pedestrian in  $\mathbf{x}$ . Thus  $y = f(\mathbf{x}; \theta)$ , where  $y$  is the predictions including the pedestrian’s candidate bounding boxes  $\{b_i | i = 1, \dots, n\}$  and confidence scores  $\{s_i | i = 1, \dots, n\}$ . The bounding box with the highest score is regarded as the final detected pedestrian.

The goal of adversarial examples  $\mathbf{x}_{adv}$  is to prevent the pedestrian detector from detecting the pedestrian. Inspired by [18], we minimize the maximum confidence score within all the bounding boxes until it falls below the detection threshold. Thus the pedestrian disappears in the thermal image. This disappear attack loss is defined as follows:

$$\mathcal{L}_{attack}(f(\mathbf{x}_{adv})) = \max_{i \in \{1, \dots, n\}} (s_i). \quad (1)$$

In this paper, we propose the adversarial infrared patch to construct the adversarial example. It can be defined as:

$$\mathbf{x}_{adv} = \mathbf{x} \odot (\mathbf{1} - \mathbf{M}) + \hat{\mathbf{x}} \odot \mathbf{M}, \quad (2)$$

where  $\odot$  is Hadamard product,  $\mathbf{M} \in \{0, 1\}^{h \times w}$  is a mask matrix used to constrain the shape and location of the infrared patches on the target object,  $\hat{\mathbf{x}} \in \mathbb{R}^{h \times w}$  denotes a cover image used to manipulate  $\mathbf{x} \in \mathbb{R}^{h \times w}$ . The infrared patches can be described as all the region where  $M_{jk} = 1$ . It is clear that the location and shape of infrared patches depend on the mask  $\mathbf{M}$ , and the content of infrared patches depends on the cover image  $\hat{\mathbf{x}}$ .

In the real world,  $\hat{\mathbf{x}}$  can not be designed freely just as the content of adversarial patches in light visible area [7, 10]. It depends on the thermal property of infrared patch material which mainly includes thermal reflectivity and thermal emissivity [20]. Since infrared patch is generally made of the same material, the values of  $\hat{\mathbf{x}}$  are usually the same. Therefore, once the material used to make infrared patches has been picked, the values of  $\hat{\mathbf{x}}$  can be determined. And then we just need to optimize the mask  $\mathbf{M}$  to generate the infrared patches to find the adversarial examples.

#### 3.2. Optimization for Mask

We here utilize a gradient-based optimization to solve for the optimal  $\mathbf{M}$ . To determine the patch’s shape and location, the solved  $\mathbf{M}$  should have the following properties: (1) The pixels with ones in  $\mathbf{M}$  should stick together, and thus they can form a valid shape on an available location. (2) The pixel values in mask  $\mathbf{M}$  should tend to one or zero as possible. However, directly solving Eq.(1) will lead to a continuous change ranging from 0 to 1 to the value in  $\mathbf{M}$ , and the pixels with high values will scatter in different location, which cannot meet the above goal. So we add the aggregation regularization and binary regularization to Eq.(1) to regularize the pixel values in the optimization, and achieve the above properties, respectively.

##### 3.2.1 Aggregation regularization

As we known, the patch areas are well-shaped, i.e., the pixels with value ones in  $\mathbf{M}$  should be clustered. To measure this aggregation of pixels in mask  $\mathbf{M}$ , we introduce local clustering coefficient [6, 21], which quantifies how close one point’s neighbours are to being a clique. For an undirected graph, the local clustering coefficient is defined as:

$$C_i = \frac{2|\{e_{jk} : v_j, v_k \in L_i\}|}{k_i(k_i - 1)}, \quad (3)$$

where  $L_i$  represents the set of immediately connected vertices with vertex  $v_i$ .  $v_j$  and  $v_k$  are two vertices that belong to  $L_i$ .  $e_{jk}$  denotes the edge connecting vertex  $v_j$  with vertex  $v_k$ ,  $k_i$  denotes the number of neighbours of vertex  $v_i$ .  $|\cdot|$  denotes the number of edges. Please refer to [6, 21] for details.

Taking Figure 3 as an example, we aim to compute the aggregation degree for the center white point  $v_i$ . Thus, there are 8 immediately connected vertices around it, which is illustrated by the blue points. According to the definition of  $e_{jk}$  in Eq.(3), we plot all the edges with red lines in Figure 3. For simplicity, we compute the edge number according to each vertex. Thus, the edge number for each vertex is given in the red text. If we define a kernel  $\mathcal{K}$  using the edge number for each vertex, i.e. the matrix in Figure 3 (b), then the edge number  $|\{e_{jk} : v_j, v_k \in L_i\}|$  in Eq.(3) can be directly computed by  $(\mathcal{K} * \mathbf{I}_{3 \times 3})/2$ , where  $\mathbf{I}_{3 \times 3}$  denotes a matrix full of one and has the same size as  $\mathcal{K}$ , and  $*$  is the convolution operation. Thus, Eq.(3) is reformulated as:

$$C_i = \frac{\mathcal{K} * \mathbf{I}_{3 \times 3}}{k_i(k_i - 1)}. \quad (4)$$

In the example of Figure 3 (a), there are 8 neighbour vertices around the center white vertex, therefore,  $k_i = 8$  here.

However, in some situations, the vertex may have a value ranging from 0 to 1, rather than 0 or 1. Therefore, we design a soft version for Eq.(4). In Figure 3, a decay factor is

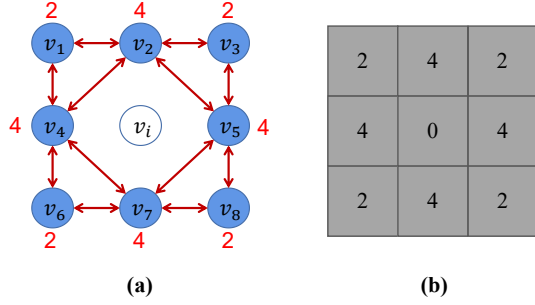


Figure 3. The left picture shows the connection among all the 8 neighbors of the center vertex. The right picture represents the kernel  $\mathcal{K}$ , where the value of the position represents the number of its immediately connected neighbours in the 8 vertices.

assigned to each vertex according to its own value and the values of its neighbour vertex. The decay factor for the vertex  $v_j$  is defined as follows:

$$\alpha_j = V_j \times \frac{1}{|L_j|} \sum_{v_k \in L_j} V_k, \quad (5)$$

where  $V_j$  is the vertex  $v_j$ 's value,  $L_j$  is the set of vertex connected with  $v_j$ , and  $V_k$  denotes the vertex's value located in  $L_j$ .  $|L_j|$  is the size of  $L_j$ . For example, for the vertex  $v_2$  in Figure.3, there are four neighbour vertex, thus  $L_2 = \{v_1, v_3, v_4, v_5\}$ , and  $|L_2| = 4$ . For the vertex  $v_1$ , there are two neighbour vertex, thus  $L_1 = \{v_2, v_4\}$ , and  $|L_1| = 2$ . After computing  $\alpha_j$  to obtain the decay factor matrix  $\mathbf{A}_{3 \times 3}$ , we replace  $\mathbf{I}_{3 \times 3}$  with  $\mathbf{A}_{3 \times 3}$  as follows:

$$C_i = \frac{\mathcal{K} * \mathbf{A}_{3 \times 3}}{k_i(k_i - 1)} = \frac{\mathcal{K} * \mathbf{A}_{3 \times 3}}{56}. \quad (6)$$

In our task, we use Eq.(6) to compute the aggregation degree  $C_{jk}$  for each element  $M_{jk}$  in  $\mathbf{M}$ . We have so far defined the  $C_{jk}$  of  $M_{jk}$ , which has nothing to do with the  $M_{jk}$  itself. To take the value of  $M_{jk}$  into consideration, we multiply the  $C_{jk}$  by  $M_{jk}$ . Finally, we take the average of all the  $C_{jk}$  into consideration, thus the aggregation regularization  $\mathcal{L}_{agg}(\mathbf{M})$  and aggregation degree  $C(\mathbf{M})$  is defined:

$$\mathcal{L}_{agg}(\mathbf{M}) = -C(\mathbf{M}) = -\frac{1}{hw} \sum_{j=1}^h \sum_{k=1}^w C_{jk} \cdot M_{jk}, \quad (7)$$

where a larger  $C(\mathbf{M})$  means a better aggregation degree.

Obviously, all the operations of the calculation of  $\mathcal{L}_{agg}$  are differentiable which can be applied in the gradient-based optimization method. Our goal is to decrease  $\mathcal{L}_{agg}$  to make the cluster of infrared patches large enough so that the points in mask  $\mathbf{M}$  gather stick on an available location.

### 3.2.2 Binary regularization

As mentioned before,  $\mathbf{M}$  is a binary matrix. To meet this goal, we use MSE (Mean Square Error) to construct the

regularization as follows:

$$\mathcal{L}_{MSE}(\mathcal{H}(\mathbf{M}), \mathbf{I}) = \frac{1}{hw} \sum_{j=1}^h \sum_{k=1}^w (\mathcal{H}(M_{jk}) - 1)^2, \quad (8)$$

where  $\mathbf{I}$  is a matrix whose each element is 1. The function  $\mathcal{H}$  here is akin to ReLU function.  $\mathcal{H}(M_{jk}) = 1$  if  $M_{jk} \leq V_{thre}$  and  $\mathcal{H}(M_{jk}) = M_{jk}$  if  $M_{jk} \geq V_{thre}$ . In the ReLU function, the derivative of point whose value is equal to zero is set as zero, while in the function  $\mathcal{H}$ , the derivative of point whose value equals to  $V_{thre}$  is also set as zero, and thus we can make  $\mathcal{H}$  differentiable like ReLU function. By minimizing Eq.(8), we only measure the distance between one and the pixels' values in  $\mathbf{M}$  which are greater than threshold  $V_{thre}$ .

In addition, because the patch should be a localized region on the target object, we add the sparse restriction on  $\mathbf{M}$ . Since  $L_0$ -norm is non-differentiable, we replace it with  $L_1$ -norm. By minimizing  $\|\mathbf{M}\|_1$ , we can select as sparse pixels as possible to make the pedestrian disappear in the image. The final regularization is defined as follows:

$$\mathcal{L}_{binary}(\mathbf{M}) = \|\mathbf{M}\|_1 + \alpha \cdot \mathcal{L}_{MSE}(\mathcal{H}(\mathbf{M}), \mathbf{I}), \quad (9)$$

where  $\alpha$  is a weighted factor. Our goal is to minimize  $\mathcal{L}_{binary}$  so that we optimize the sparse and binary  $\mathbf{M}$ .

### 3.2.3 Total loss function

Now, we combine three losses mentioned above as our total loss, and keep balance among the losses with weighted factors  $\lambda_1$  and  $\lambda_2$ . The final objective function is given by:

$$\mathcal{L}_{obj} = \mathcal{L}_{attack}(f(\mathbf{x}_{adv})) + \lambda_1 \mathcal{L}_{binary}(\mathbf{M}) + \lambda_2 \mathcal{L}_{agg}(\mathbf{M}) \quad (10)$$

$$\mathbf{M}^* = \arg \min_{\mathbf{M}} \mathcal{L}_{obj}(f(\mathbf{x}_{adv})), \quad (11)$$

where the weighted factors can be obtained by experience to meet our requirements. In the implementation, to ensure the infrared patches to attach on the target object, they should be restricted to the region of the target object in the infrared image. In detail, given an available binary  $\mathbf{M}_{obj}$  as the object mask to locate the target object, we can process  $\mathbf{M} = \mathbf{M} \cdot \mathbf{M}_{obj}$  in Eq.(2) to perform the following optimization.  $\mathbf{M}_{obj}$  can be obtained via object detection or segmentation.

### 3.2.4 Optimization

To solve Eq.(10), we use the gradient descent algorithm with momentum [11]. Let  $t$  denote the  $t$ -th iteration,  $\mathbf{x}_t^{adv}$ ,  $\mathbf{M}_t$  denote the adversarial image and infrared mask respectively in the  $t$ -th iteration, then:

$$\mathbf{x}_{t+1}^{adv} = \mathbf{x} \odot (1 - \mathbf{M}_{t+1}) + \hat{\mathbf{x}} \odot \mathbf{M}_{t+1}, \quad (12)$$

$$\mathbf{M}_{t+1} = (\mathbf{M}_t - \epsilon \cdot g_{t+1}) \odot \mathbf{M}_{obj}, \quad (13)$$

$$g_{t+1} = \mu \cdot g_t + \frac{\nabla_{\mathbf{M}} \mathcal{L}_{obj}(f(\mathbf{x}_{adv}), y; \mathbf{M}_t)}{\|\nabla_{\mathbf{M}} \mathcal{L}_{obj}(f(\mathbf{x}_{adv}), y; \mathbf{M}_t)\|_1}. \quad (14)$$

In practice, however, we find some elements in  $\mathbf{M}_{t+1}$  output by Eq.(14) are still scattered. To smooth these elements, we fine-tune the gradient in the optimization process as follows:

$$g_{t+1} = g_{t+1} \odot \text{Norm}(\mathbf{M}_t * \mathcal{K}_{gau}), \quad (15)$$

where  $\text{Norm}(\cdot)$  is normalization,  $\mathcal{K}_{gau}$  is Gaussian kernel,  $*$  is convolution. Eq.(15) considers the elements' context information to remove outlier points, making the shape smooth.

### 3.3. Analysis for Cover Image

The infrared thermal imagery of an object is determined by the thermal radiation emitted by the object. Based on this feature, we put a patch of thermal insulation material at a specific location of the target to reduce the thermal radiation, thus changing the distribution of the infrared radiation of the object to fool the detector.




Figure 4. The comparison of infrared images taken by the infrared camera between (b) normal object and (c) object with patches made by the (a) aerogel material.

In practice, we choose the aerogel material as the thermal insulation material for its low thermal radiation and low thermal emissivity. The material and its heat-insulating effect can be seen in Figure 4. We can see the material on the object can shield most of the thermal radiation of the target object and reduce the thermal radiation of the corresponding area to a certain extent. We find that the gray values corresponding to the region with material in the infrared image are almost the same, so these gray values are set to  $\hat{\mathbf{x}}$ .

### 3.4. The Overall Algorithm

The overall algorithm for generating adversarial infrared patches is summarized in Algorithm 1. Given the pedestrian detector and infrared image, Algorithm 1 will output the adversarial infrared patch and adversarial image. The process of optimized mask  $\mathbf{M}$  is given in Figure 5. We see that the mask gradually converges to a stable shape and location.

## 4. Experiments

In this section, we first perform the ablation study for our method in the digital world against pedestrian detector, and then give the attack performance in the physical world

---

### Algorithm 1 Optimization for Adversarial Infrared Patches

---

**Input:** Pedestrian detector  $f(\cdot)$ , clean infrared image  $\mathbf{x}$ , cover image  $\hat{\mathbf{x}}$ , restriction region  $\mathbf{M}_{obj}$ , max iteration  $T$ , max size  $\varepsilon_{max}$ , threshold  $s_{thr}$ .

**Output:** Infrared patch  $\mathbf{M}^*$ , and adversarial example  $\mathbf{x}_{adv}^*$

- 1: Initialize  $\mathbf{M}_0$  randomly in  $[0, 1]$ ,  $g_0 = 0$ .
  - 2: Restrict  $\mathbf{M}_0$  in the object region:  $\mathbf{M}_0 = \mathbf{M}_0 \odot \mathbf{M}_{obj}$
  - 3: **for**  $t = 0$  to  $T$  **do**
  - 4:   Compute the gradient  $g_{t+1}$  according to Eq.(14)
  - 5:   Finetune the gradient  $g_{t+1}$  according to Eq.(15)
  - 6:   Update the mask  $\mathbf{M}_{t+1}$  according to Eq.(13)
  - 7:   Compute  $\mathbf{x}_{t+1}^{adv}$  according to Eq.(12)
  - 8:   Perform the prediction  $y = f(\mathbf{x}_{t+1}^{adv}; \theta)$ , and select the top-1 bounding box  $(b_1, s_1)$
  - 9:   **if**  $s_1 \leq s_{thr}$  and  $\|\mathbf{M}_{t+1}\|_1 \leq \varepsilon_{max}$  **then**
  - 10:      $\mathbf{M}^* = \mathbf{M}_{t+1}$ ;  $\mathbf{x}_{adv}^* = \mathbf{x}_{t+1}^{adv}$ ; **break**;
  - 11:   **end if**
  - 12: **end for**
  - 13: **return**  $\mathbf{M}^*$ ,  $\mathbf{x}_{adv}^*$
- 

against pedestrian detector and vehicle detector, respectively. Finally, the result against defense is given. More experiments can be found in the *Supplemental Material*.

### 4.1. Simulation of Physical Attacks

**Dataset:** We use the FLIR ADAS dataset<sup>1</sup> to simulate the physical attacks. Similar to [29] and [30], we select the images containing pedestrian as our dataset. The training set includes 7873 images, and the test set includes 2027 images. We then choose the test images that can be successfully detected by the target model with a high probability as the final images to perform attacks. Thus, the clean AP is 100%. **Target detector:** For pedestrian detection task, we also choose YOLOv3 [14] based detector for its high speed, which is the same to [29] and [30]. We choose the officially pretrained weights as the initialized weights and then re-train the model on the FLIR ADAS training dataset.

**Metrics:** Attack Success Rate (ASR) and Average Precision (AP) are used to evaluate the attack performance. ASR denotes the ratio of successfully attacked images out of all the test images. AP is computed by measuring the region under the Precision-Recall (PR) curve.

#### 4.1.1 Comparisons with SOTA methods

We first test the attack performance of our method in the digital world, and compare with adversarial clothing [29] and adversarial bulbs [30] in the same dataset. The Precision-Recall curves are plot in Figure 6 (a) and the corresponding AP values are also listed, where we see that the AP of pedestrian detector drops to 12.05% from 100% after the attack

<sup>1</sup><https://www.flir.com/oem/adas/adas-dataset-form/>




Figure 5. The optimization process of mask  $M$  for every 5 iterations, where we can see the mask gradually becomes aggregated and stable.




Figure 6. Comparisons with SOTA infrared attacks in digital world.

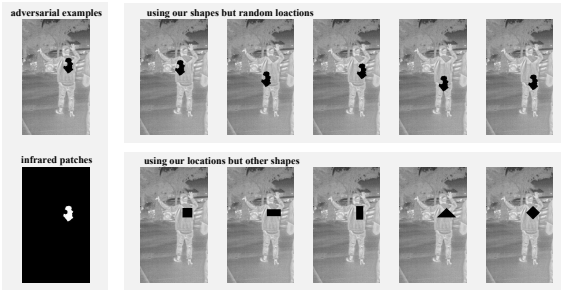


Figure 7. The experimental setting of ablation study for the infrared patches' shapes and locations.

by our method, which is obviously lower than adversarial bulbs' 35.88%, and slightly lower than adversarial clothing's 12.30%. These results verify the advantage of our method. A qualitative example is given in Figure 6 (b).

Table 1. Ablation study for infrared patches' shapes and locations.

Our location	Our shape	ASR
—	—	0%
✓	—	54.16%
—	✓	59.68%
✓	✓	<b>93.60%</b>

#### 4.1.2 Effects of shapes and locations

To investigate the effects of different shapes and locations, we here give the ablation study. Specifically, we first generate adversarial infrared patches by solving Eq.(10) to achieve the optimal shapes and locations (the left sub-figure in Figure 7) for a given pedestrian. Then, we conduct two additional experiments. One is fixing the optimal shapes while using random locations (the top row in the right sub-figure in Figure 7), then computing the average ASR under these five random locations. This setting tests the effects of optimal locations to the attacks. The other is fixing the optimal locations while changing different patches' shapes (the bottom




Figure 8. Example for infrared mask  $M$ . (a) With only  $\mathcal{L}_{attack}$ . (b) With  $\mathcal{L}_{attack}$  and  $\mathcal{L}_{binary}$ . (c) With  $\mathcal{L}_{attack}$ ,  $\mathcal{L}_{binary}$  and  $\mathcal{L}_{agg}$ .

row in the right sub-figure in Figure 7). These shapes can be square, horizontal rectangle, vertical rectangle, triangle, rhombus, etc. Then we compute the average ASR under these five shapes. This setting tests the effects of optimal shapes to the attacks. The results are listed in Table 1, where we can see (1) With the optimal shapes and locations, adversarial infrared patches achieve the best ASR (93.60%). (2) When only using the optimal shapes or locations, there still exists the attacking ability for adversarial infrared patches, but the performance is weaker. And the locations and shapes have a similar attacking ability (54.16% vs 59.68%). (3) When neither our location nor our shape is used, the ASR decreases to 0%. The contrast in Table 1 demonstrates the importance of patches' shapes and locations to the attack.

Table 2. Ablation study for different loss functions.

Loss functions	ASR	Aggregation
$\mathcal{L}_{attack}$	100%	0.0057
$\mathcal{L}_{attack} + \mathcal{L}_{binary}$	99.60%	0.0427
$\mathcal{L}_{attack} + \mathcal{L}_{binary} + \mathcal{L}_{agg}$	93.60%	0.4651

#### 4.1.3 Effects of loss functions

To visualize the impact of different loss functions, we stepwise add  $\mathcal{L}_{attack}$ ,  $\mathcal{L}_{binary}$ , and  $\mathcal{L}_{agg}$  to total loss as Eq.(10) to optimize the infrared mask  $M$  and record the mask  $M$  after the optimization. An example of generated mask  $M$  on a pedestrian is given in Figure 8. We can see that optimized with only  $\mathcal{L}_{attack}$ , the infrared mask  $M$  presents "QR code" texture. When binary regularization  $\mathcal{L}_{binary}$  is added, the elements in infrared mask  $M$  look like grainy dots, and more regions turn to black areas. It proves that the  $\mathcal{L}_{binary}$  can effectively make the values of mask  $M$  get close to zero or one. Finally, when aggregation regularization  $\mathcal{L}_{agg}$  is added, the grainy elements immediately gather together, and the infrared patch becomes well-shaped on a specific location, which can be found that it is easy to be implemented in the real world. Besides, we also give the quantitative results




Figure 9. Visual examples of physical attacks with infrared patches under various angles, distances, postures, and scenes.

Table 3. ASR in the physical world when changing distances, postures, angles, and scenes of pedestrian captured by infrared cameras.

	0°	±10°	±20°	±30°	dist.	pos.	outdoor
ASR	94.67%	92.04%	89.05%	83.38%	77.25%	76.70%	87.70%




Figure 10. Visualization results of physical attacks with adversarial infrared patches. The top row lists the detection results on clean infrared images. And the bottom row lists the detection results on adversarial infrared images. We consider the front view in (a), different shooting angles in (b), sitting posture in (c), different shooting distances in (d), and outdoor scene in (e).

of different loss functions in Table 2, where we see with the integration of  $\mathcal{L}_{binary}$  and  $\mathcal{L}_{agg}$ , the ASR has a slight reduction from 100% to 93.60%, but the aggregation degree shows a remarkable increase. A larger aggregation degree means a good physical implementation.

## 4.2. Attacks in the Physical World

### 4.2.1 Physical attacks against pedestrian detector

To test the performance of infrared patches under different physical conditions, we design various situations to conduct physical attacks and record the video to calculate the ASR. In default, we take videos at the 4 meters away from the standing pedestrian in the frontal view ( $0^\circ$ ) within the indoor scene. For the angle problem, we change the angles with  $\pm 10^\circ$ ,  $\pm 20^\circ$ ,  $\pm 30^\circ$ . For distance problem, we move the camera to 6 meters from the default 4 meters. For posture problem, we change the pedestrian’s posture from the default standing to sitting pose. For the scene problem, we change to outdoor from the default indoor. The visual examples of these situations are listed in Figure 9. During the shooting process, the pedestrian is asked to move their bodies within the range of  $5^\circ$  of the current posture to take videos. For each situation, we shoot for 20 seconds at ten frames a second (about 200 frames in total) and compute the ASR versus the captured video. The pedestrian detection threshold is set as

0.5. Table 3 lists the quantitative results. It can be seen that our infrared patches achieve a high ASR (94.67%) in the frontal view. When the shooting angle changes, the ASR still maintains a high value (92.04%, 89.05%, and 83.38%). When changing the distance from 4 meters to 6 meters, the ASR decreases to 77.25%. When changing the posture from standing to sitting, the ASR decreases to 76.70%, and ASR decreases to 87.70% when the scene is changed to outdoor. These results show that the impact of different shooting situations is relatively small to the infrared patches. In other words, as long as the shape of our infrared patch on the object can be completely captured by the camera, the effects of adversarial attacks can be maintained. Some qualitative results before and after attaching the infrared patch on the pedestrian in different settings are given in Figure 10.

We compare our infrared patches with the SOTA physical infrared attack: adversarial clothing [29]<sup>2</sup>. We use the data provided in [29] and conduct our physical attack. For shooting angles, we select  $0^\circ$ ,  $\pm 10^\circ$ ,  $\pm 20^\circ$ , and  $\pm 30^\circ$  and fix the distance of 5 meters. For shooting distance, we select 3 meters, 4 meters, 5 meters, and 6 meters. We set the detection threshold as 0.7 like [29]. Finally, we compare the time cost for infrared clothing and infrared patches in real world to evaluate their implementation difficulty of physical attacks. The above results are in Figure 11. From the results, we can see that: (1) Our adversarial infrared patches have better performance than adversarial clothing when the pedestrian is captured in the frontal view, and show the competitive performance when the angle changes far away from the frontal view. Because our method can maintain the good attacks within  $60^\circ$ , it is potential to extend it to achieve  $360^\circ$  attack by attaching at least six infrared patches around the pedestrian. (2) Our infrared patches have a similar performance with infrared clothing when the distance is located between the 3-th meter and 6-th meter, which shows the robustness against different distances. (3) The production time cost for infrared patches in real world (only about 0.5 hours) is far

<sup>2</sup>We don’t compare with adversarial bulbs [30] because they don’t report corresponding physical results under different angles and distances, and reproducing the physical attacks is difficult owing to its poor implementation.

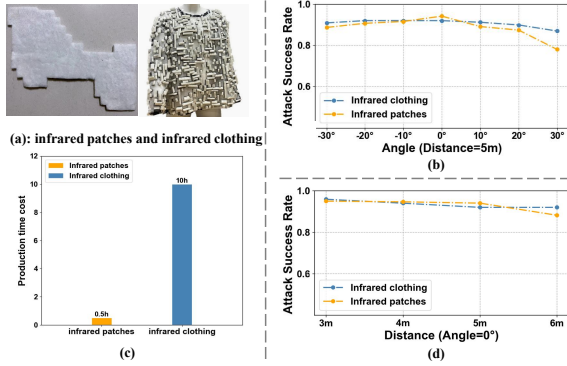


Figure 11. The comparisons of ASRs in real world at different angles (b), different distances (d). We also report the production time cost (c) of generating a physical adversarial example in real world between infrared clothing and infrared patches (a).

less than the infrared clothing (about 10 hours), which shows the production process of our infrared patches is much more convenient than infrared clothing. All in all, our adversarial infrared patches achieve the competitive attack performance in various situations but only cost five percent of their time to construct the physical attack compared with adversarial clothing attack, which demonstrates our advantage.

#### 4.2.2 Physical attacks against vehicle detector

Besides the pedestrian detection task, our adversarial infrared patch can also be applied to other tasks. In this section, we conduct the experiments of physical attacks against vehicle detection in the aerial images [13]. For that, a two-stage object detector: Rotated Cascade R-CNN [31] is chosen as the target detector. In this way, we can verify our adversarial infrared patch against not only the one-stage YOLO detector above but also the two-stage R-CNN detector. Meanwhile, the experiments can also verify the wide application in both the horizontal object detector and rotated object detector. Like the pedestrian detection task, we record the vehicle video captured in different angles and distances, and then compute the ratio of successfully attacked video frames out of all the frames as the ASR. Thus, we obtain a video with 750 frames and 696 frames are successfully attacked, leading to a 92.8% ASR versus the vehicle detection.

The results are given in Figure 12, where the left denotes the qualitative examples and the right denotes the quantitative ASR against different angles and distances. To give a clear comparison, we list two vehicles in an image, where the vehicle attacked by the infrared patch cannot be detected but the clean vehicle is successfully detected. Additionally, we see the patch' size is small versus the vehicle, but it indeed performs the successful adversarial attack, which can show the vulnerability of vehicle detector against the infrared patch. The right quantitative results show the high attack performance versus different angles and distances. It




Figure 12. The qualitative and quantitative results of attacking the vehicle detector in the physical world.

shows the similar trends with the pedestrian detection task. Because there are no published physical attacks against the vehicle detection task, we don't give the comparison here.

#### 4.3. Defenses against Adversarial Infrared Patches

We test three typical methods to defend our attack method in the digital world. One is pre-processing defense: spatial smoothing [27], the second one is adversarial training [5], the third one is to randomize gradients [12]. Results are given in Table 4, where we see that (1) after spatial smoothing, ASR only drops 8%. This is reasonable because the cover image in our method has the same value, the smooth operation cannot change the distribution. (2) After randomizing gradients by adding random noises on the image, ASR is 14.8% lower than no defense. (3) After adversarial training, ASR drops 29%, which is still acceptable. It shows the robustness of our adversarial infrared patches to defense methods.

Table 4. Results against the defense methods.

Defenses	No Defense	[27]	[12]	[5]
ASR	93.6%	85.6%	78.8%	64.6%

### 5. Conclusion

In this paper, we proposed physically feasible adversarial infrared patches with learnable shapes and locations, which can be solved through a carefully designed gradient-based optimization method. We constructed the infrared patches with thermal insulation material. The production of infrared patches was simple and the physical attack can be easy to realize. It only needed 0.5 hours to construct the adversarial patch in the physical world, which significantly outperformed the SOTA infrared attack methods. Experiments on the pedestrian detection and vehicle detection verified the effectiveness of the proposed method.

### Acknowledgment

This work is supported by National Key R&D Program of China (No.2020AAA0104002), National Natural Science Foundation of China (No.62076018).



## References

- [1] Tom B Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. *arXiv preprint arXiv:1712.09665*, 2017.
- [2] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.
- [3] DeMarcus Edwards and Danda B Rawat. Study of adversarial machine learning with infrared examples for surveillance applications. *Electronics*, 9(8):1284, 2020.
- [4] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust physical-world attacks on deep learning visual classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1625–1634, 2018.
- [5] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [6] Paul W Holland and Samuel Leinhardt. Transitivity in structural models of small groups. *Comparative group studies*, 2(2):107–124, 1971.
- [7] Stepan Komkov and Aleksandr Petiushko. Advhat: Real-world adversarial attack on arface face id system. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 819–826. IEEE, 2021.
- [8] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- [9] Uche M Osahor and Nasser M Nasrabadi. Deep adversarial attack on target detection systems. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, volume 11006, page 110061Q. International Society for Optics and Photonics, 2019.
- [10] Mikhail Pautov, Grigorii Melnikov, Edgar Kaziakhmedov, Klim Kireev, and Aleksandr Petiushko. On adversarial patches: real-world attack on arface-100 face recognition system. In *2019 International Multi-Conference on Engineering, Computer and Information Sciences (SIBIRCON)*, pages 0391–0396. IEEE, 2019.
- [11] Ning Qian. On the momentum term in gradient descent learning algorithms. *Neural networks*, 12(1):145–151, 1999.
- [12] Zeyu Qin, Yanbo Fan, Hongyuan Zha, and Baoyuan Wu. Random noise defense against query-based black-box attacks. *Advances in Neural Information Processing Systems*, 34:7650–7663, 2021.
- [13] Sebastien Razakarivony and Frederic Jurie. Vehicle detection in aerial imagery: A small target detection benchmark. *Journal of Visual Communication and Image Representation*, 34:187–203, 2016.
- [14] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- [15] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [16] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [17] Mahmood Sharif, Sruti Bhagavatula, Lujio Bauer, and Michael 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*, pages 1528–1540, 2016.
- [18] Dawn Song, Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Florian Tramer, Atul Prakash, and Tadayoshi Kohno. Physical adversarial examples for object detectors. In *12th USENIX workshop on offensive technologies (WOOT 18)*, 2018.
- [19] Frédéric Suard, Alain Rakotomamonjy, Abdelaziz Benschraï, and Alberto Broggi. Pedestrian detection using infrared images and histograms of oriented gradients. In *2006 IEEE Intelligent Vehicles Symposium*, pages 206–212. IEEE, 2006.
- [20] Michael Vollmer. Infrared thermal imaging. In *Computer Vision: A Reference Guide*, pages 666–670. Springer, 2021.
- [21] Duncan J Watts and Steven H Strogatz. Collective dynamics of ‘small-world’ networks. *nature*, 393(6684):440–442, 1998.
- [22] Xingxing Wei, Ying Guo, and Jie Yu. Adversarial sticker: A stealthy attack method in the physical world. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [23] Xingxing Wei, Ying Guo, Jie Yu, and Bo Zhang. Simultaneously optimizing perturbations and positions for black-box adversarial patch attacks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [24] Xingxing Wei, Bangzheng Pu, Jiefan Lu, and Baoyuan Wu. Physically adversarial attacks and defenses in computer vision: A survey. *arXiv preprint arXiv:2211.01671*, 2022.
- [25] Qihao Weng. Thermal infrared remote sensing for urban climate and environmental studies: Methods, applications, and trends. *ISPRS Journal of photogrammetry and remote sensing*, 64(4):335–344, 2009.
- [26] Chaowei Xiao, Bo Li, Jun-Yan Zhu, Warren He, Mingyan Liu, and Dawn Song. Generating adversarial examples with adversarial networks. *arXiv preprint arXiv:1801.02610*, 2018.
- [27] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155*, 2017.
- [28] Chenglin Yang, Adam Kortylewski, Cihang Xie, Yinzi Cao, and Alan Yuille. Patchattack: A black-box texture-based attack with reinforcement learning. In *European Conference on Computer Vision*, pages 681–698. Springer, 2020.
- [29] Xiaopei Zhu, Zhanhao Hu, Siyuan Huang, Jianmin Li, and Xiaolin Hu. Infrared invisible clothing: Hiding from infrared detectors at multiple angles in real world. *arXiv preprint arXiv:2205.05909*, 2022.
- [30] Xiaopei Zhu, Xiao Li, Jianmin Li, Zheyao Wang, and Xiaolin Hu. Fooling thermal infrared pedestrian detectors in real world using small bulbs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3616–3624, 2021.
- [31] Yixing Zhu, Chixiang Ma, and Jun Du. Rotated cascade r-cnn: A shape robust detector with coordinate regression. *Pattern recognition*, 96:106964, 2019.