

Unified Adversarial Patch for Cross-modal Attacks in the Physical World

Xingxing Wei^{1,2,*}, Yao Huang¹, Yitong Sun¹, Jie Yu³

¹ Institute of Artificial Intelligence, Beihang University, Beijing 100191, China

² Hangzhou Innovation Institute, Beihang University, Hangzhou 311228, China

³ School of Computer Science and Engineering, Beihang University, Beijing 100191, China

{xxwei, y.huang, yt.sun, sy2106137}@buaa.edu.cn

Abstract

Recently, physical adversarial attacks have been presented to evade DNNs-based object detectors. To ensure the security, many scenarios are simultaneously deployed with visible sensors and infrared sensors, leading to the failures of these single-modal physical attacks. To show the potential risks under such scenes, we propose a unified adversarial patch to perform cross-modal physical attacks, i.e., fooling visible and infrared object detectors at the same time via a single patch. Considering different imaging mechanisms of visible and infrared sensors, our work focuses on modeling the shapes of adversarial patches, which can be captured in different modalities when they change. To this end, we design a novel boundary-limited shape optimization to achieve the compact and smooth shapes, and thus they can be easily implemented in the physical world. In addition, to balance the fooling degree between visible detector and infrared detector during the optimization process, we propose a score-aware iterative evaluation, which can guide the adversarial patch to iteratively reduce the predicted scores of the multi-modal sensors. We finally test our method against the one-stage detector: YOLOv3 and the two-stage detector: Faster RCNN. Results show that our unified patch achieves an Attack Success Rate (ASR) of 73.33% and 69.17%, respectively. More importantly, we verify the effective attacks in the physical world when visible and infrared sensors shoot the objects under various settings like different angles, distances, postures, and scenes.

1. Introduction

Deep Neural Networks (DNNs) are vulnerable to adversarial examples [18], which pose a serious security threat to DNNs-based object detectors in the physical world [31, 34, 33]. These physical attacks can help to evaluate the

*Corresponding author

Table 1. Various adversarial attacks in different settings.

	Digital world	Physical world
Single-modal	[28],[4],[30],[8], [11],[13] etc.	[20], [17] [10],[6] [26],[29],[33] etc.
Cross-modal	[1],[21] etc.	Ours

robustness of DNNs deployed in real-life systems, which have important practical values. Recently, RGB sensors and thermal infrared sensors have been simultaneously used in many safety-critical tasks such as security monitoring. When performing object detection tasks, visible images could provide abundant information on the target’s texture in the daytime, while infrared images could display the target’s thermal distribution at night. Thus, combining visible images and infrared images together will result in a round-the-clock application. To evade object detectors under such multi-modal imaging scenarios, a necessary way is to develop a cross-modal physical attack to fool visible object detectors and infrared object detectors at the same time.

However, current physical attacks are limited in the single-modal domain. For example, some studies [20, 6, 27] hide from object detectors in the visible modality, and some studies [34, 33, 25] hide from object detectors in the infrared modality. Because visible sensors and infrared sensors have different imaging mechanisms, these single-modal physical attacks cannot simultaneously attack multi-modal object detectors. Specifically, the generated perturbations in the visible modality can not be captured by infrared sensors, and in turn, changing the object’s thermal radiation would not be reflected in the visible light domain. There also exist some cross-modal attacks in the digital world [1, 21, 24], but they usually aim to modify the image pixels or point clouds after sensors’ imaging, ignoring the different imaging mechanisms across sensors, and thus cannot transfer to physical world well.

Based on these discussions, this paper aims to design a **unified cross-modal attack in the physical world** to fill in the gap of this area, as illustrated in Table 1. Generally speaking, the lack of features that can operate in different

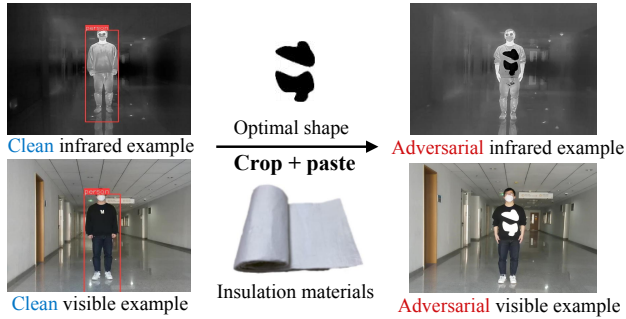


Figure 1. The generation process for our unified cross-modal adversarial patches. We see the pedestrian cannot be detected after the patches are pasted on the pedestrian in the physical world.

modalities is the reason to limit the availability of cross-modal physical attacks. Inspired by [5] and [25], we utilize adversarial patches [2] and further optimize their shapes to perform attacks. Because patches’ shapes are a universal attribute for different imaging mechanisms, their changes can be well captured by both visible sensors and infrared sensors, making it an appropriate form to perform cross-modal attacks. However, there still exist two challenges to meet our goal: (1) The current shape models are either heteromorphic [5] or change-limited [25], leading to the difficulty of physical implementation or poor attacks because of finite searching space. So how to design a flexible shape optimization method while keeping shapes’ smoothness is the first challenge. (2) During the optimization for the unified physical attacks, the drop speed of visible object detector and infrared object detector may be inconsistent, which will result in the successful evasion for one object detector but failed evasion for another object detector. Therefore, how to balance the performance of two object detectors from different domains is another challenge.

To meet the above challenges, we design a novel boundary-limited shape optimization method to achieve the compact and smooth shapes, and thus they can be easily implemented in the physical world. Moreover, the shapes are flexible, and can provide a huge searching space to find the optimal shape to achieve a successful cross-modal attack. In addition, to balance the fooling degree between visible detector and infrared detector during the optimization process, we propose a score-aware iterative evaluation, which can guide the adversarial patch to iteratively reduce the predicted scores of the multi-modal sensors. When applying to physical implements, we only need to print the simulated results in the digital world and crop them with insulation materials for patches. An example of our cross-modal adversarial patch against visible and infrared pedestrian detectors is shown in Figure 1.

The contributions of this paper are as follows:

- We propose a unified adversarial patch to perform cross-modal attacks. To the best of our knowledge, it is

the first work to simultaneously evade visible detectors and infrared detectors in the physical world.

- We design two novel techniques: boundary-limited shape optimization and score-aware iterative evaluation, to achieve the feasible patches in the digital world while balancing the multi-modal detectors.
- We verify our cross-modal patches in the pedestrian detection task not only in the digital world but also in the physical world. Experimental results show that cross-modal patches can work well in various angles, distances, postures, and scenes.

2. Related Works

2.1. Attacks in the physical world

Since Kurakin *et al.* [11] verify the feasibility of adversarial attacks in the physical world, many physical attacks have been proposed for more application value, like Sharif *et al.* [16]’s adversarial glasses in face recognition systems, Eykholt *et al.* [7]’s adversarial graffiti in automatic driving tasks. Among them, adversarial patch [3] is one of the mainstream physical attacks. Different from Lp-norm-based adversarial perturbations, the adversarial patch doesn’t need to restrict the perturbations’ magnitude, which is always in the form of a patch and thus more suitable to the physical world. Like Xu *et al.* [31], they design a special adversarial T-shirt, which can apply the deformable adversarial patch on the T-shirts to fool the person detector. When referring to infrared detectors, Zhu *et al.* propose adversarial bulbs [34] based on the extra heat source and invisible cloth [33] based on QR code pattern that could successfully attack infrared detectors, but compared with adversarial patches, they are complicated to implement in the physical world and lose effects in the visible modality.

2.2. Shape model in adversarial attacks

In addition to generating perturbations for the traditional adversarial patch, some studies explore the patches’ shapes to adversarial attacks. For example, Chen *et al.* [5] propose a deformable patch representation, but their deformation shapes mainly rely on an central point and rays, which causes the shape to be too sharp and not natural enough. In the meanwhile, their work still needs to optimize the contents on the patch to some extent and mainly focus on tasks of image classification. Besides, Wei *et al.* [25] propose a hotcold block based on Warming Paste and Colding Paste to attack infrared detectors, but their deformation of patches is limited by the manually set nine-square-grid states, which greatly decreases the search space of patches’ shapes.

Overall, none of the above methods can simultaneously attack the two different modalities. Our work will combine

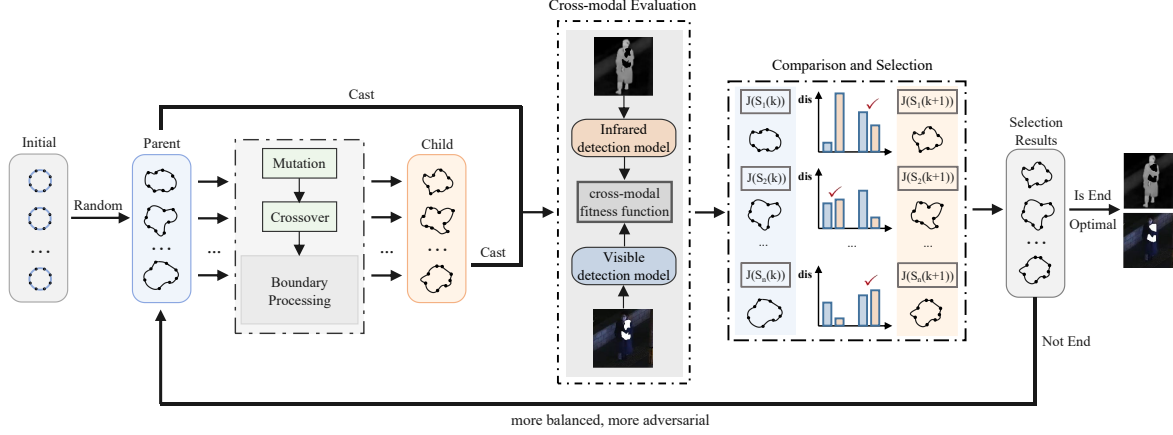


Figure 2. An overview of cross-modal adversarial patches' generation based on the Differential Evolution (DE) framework. The initial population is a series of circles. Then, combining mutation, crossover and boundary processing, a child population made up of natural shapes is generated. With a special cross-modal evaluation, we compare the parent population with the child population and select better individuals, encouraging the population to become more balanced and more adversarial. Finally, the optimal individual will be printed.

adversarial patches with cross-modal attacks from the angle of a natural shape optimization.

3. Methodology

In this section, we choose the pedestrian detection as the target task to introduce the details of our method.

3.1. Problem Formulation

In the pedestrian detection task, given a clean visible image x_{vis} and a clean infrared image x_{inf} , the goal of a unified cross-modal adversarial attack is to make the visible detector $f_{vis}(\cdot)$ and infrared detector $f_{inf}(\cdot)$ simultaneously unable to detect the pedestrian in the perturbed visible image x_{vis}^{adv} and perturbed infrared image x_{inf}^{adv} . The formulation can be expressed as follows:

$$\max(f_{vis}(x_{vis}^{adv}), f_{inf}(x_{inf}^{adv})) < thre \quad (1)$$

where $f_{vis}(x_{vis}^{adv})$ and $f_{inf}(x_{inf}^{adv})$ represent the confidence scores of detected pedestrians in the visible modality and infrared modality, and $thre$ is a pre-defined threshold.

The perturbed visible image x_{vis}^{adv} and the perturbed infrared image x_{inf}^{adv} with the unified adversarial patch can be generated as Eq.(2) and Eq.(3).

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

$$x_{inf}^{adv} = x_{inf} \odot (1 - M) + \hat{x}_{inf} \odot M, \quad (3)$$

where \odot is Hadamard product, $M \in \{0, 1\}^{h \times w}$ is a mask matrix used to constrain the shape and location of the cross-modal patches on the target object, $\hat{x}_{vis} \in R^{h \times w}$ denotes a cover image used to manipulate x_{vis}^{adv} , $\hat{x}_{inf} \in R^{h \times w}$ denotes a cover image used to manipulate x_{inf}^{adv} . The values of

these two cover images are obtained by the visible sensor and infrared sensor versus the unified insulation material in the physical world. The unified patches can be described as all the regions where $M_{ij} = 1$.

In the real application, we use the aerogel material to implement our unified adversarial patch. When infrared sensor shoots, it will show good insulation effects, and thus can change the thermal distribution of the target pedestrian. When visible sensor shoots, its color is white and can show a difference with the pedestrian. Based on this, our method optimizes a cross-modal M to learn an adversarial shape, finally working well in both modalities.

3.2. Shape Representation

To model a shape, we need to determine the shape representation. For that, we first define some points as the basic elements. Then we connect these anchor points to construct the contour, representing our shape. Here are the details.

Multi-anchor Representation: Unlike Chen *et al.* [5] using a central point and corresponding rays to form a polygon contour, we only use points, distributing multiple anchor points on the patch contour. Then, we can directly adapt the patch contour's shape by changing the coordinates of the points. Owing to the design of multi-anchor representation, we can have more flexible shape variations and a broader search space.

We illustrate this process in Figure 3 (a), where points in the dotted line denote the initialized location. The arrows show the changed direction for some points. We can adjust the coordinates of the points to control the movement of each point and thus, the direction of our points' movement will not be limited to one single direction.

Smooth Spline Connection: To ensure shapes' natural-

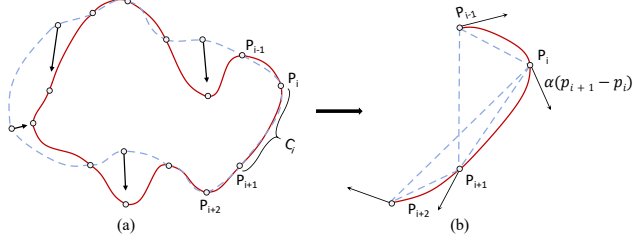


Figure 3. Subfigure (a) shows the process of moving the anchor points to change the shape, where the dotted line denotes the initialized location, and arrows show the changed direction. Subfigure (b) is an example of the curve segment C_i 's connection via Catmull-Rom spline interpolation.

ness, we use centripetal Catmull–Rom spline [22] to connect such anchor points. The process of spline can be formulated as follows:

As shown in Figure 3 (b), to generate a curve segment C_i between P_i and P_{i+1} , we will use four anchor points $P_{i-1}, P_i, P_{i+1}, P_{i+2}$ and a centripetal Catmull-Rom spline function $CCRS(\cdot)$ [22]. With the centripetal Catmull-Rom spline function, first, we can ensure that a loop or self-intersection doesn't exist within a curve segment. Second, a cusp will never occur within a curve segment. Finally, it can also follow the anchor points more tightly. C_i can be formulated as follows:

$$C_i = CCRS(P_{i-1}, P_i, P_{i+1}, P_{i+2}) \quad (4)$$

When n curve segments C_0, C_1, \dots, C_{n-1} are combined, the patch contour M_{con} can be written as:

$$M_{con} = \{C_i | 0 \leq i \leq n - 1\} \quad (5)$$

With M_{con} being closed, we can easily obtain M :

$$M(x) = \begin{cases} 1, & x \text{ inside } M_{con} \\ 0, & x \text{ out of } M_{con} \end{cases} \quad (6)$$

After Eq.(6), we can represent the patch's shape. The detailed formula for this process can be found in *Supplemental Material*.

3.3. Shape Optimization for Cross-modal Attacks

In practice, the details of $f_{vis}(\cdot)$ and $f_{inf}(\cdot)$ are usually unknown to the adversary, so it is intractable to optimize the anchor points via a gradient-based optimization method. Considering this background, we carry out a score-based black-box attack by querying the object detector to obtain the confidence scores of detected pedestrians. From the above, we formulate the whole unified cross-modal attacks into the Differential Evolution (DE) framework.

3.3.1 Formulation Overview Using DE

The Differential Evolution (DE) consists of four parts: starting from an initial population, using the crossover and mutation to generate the offspring population, making the fittest survive according to the fitness function, and finding the appropriate solution in the iterative evolution process.

In our case, a population represents the anchor points $\{P_j | j = 1, \dots, n\}$. Given the population size Q , the k -th generation solutions $S(k)$ is represented as:

$$S(k) = \{S_i(k) | \theta_j^L \leq S_{ij}(k) \leq \theta_j^U, 1 \leq i \leq Q, 1 \leq j \leq n\} \quad (7)$$

where $S_i(k)$ is the i -th patch's shape, and $S_{ij}(k)$ represents the j -th anchor point of $S_i(k)$ in the k -th generation. θ_j^L and θ_j^U together make up the feasible region B_j , which is the moving range of the j -th anchor point in each patch shape.

In the $k + 1$ generation of DE, the solution $S(k + 1)$ is achieved via crossover, mutation, and selection based on $S(k)$. A fitness function is applied on $S_i(k)$ to evaluate its attack effectiveness. During this process, the fitness function only utilizes the confidence scores of object detectors. We will give the detailed definition of feasible region in Section 3.3.2, and fitness function in Section 3.3.3.

3.3.2 Boundary-limited Shape Optimization

Obviously, if there is no boundary restriction for the anchor points during the crossover and mutation, some bad situations will occur during shape optimization, as shown in Figure 4 (d, e, f), such as boundary line crossings (d), anchor points stuck in narrow places (e), and anchor points outside the human body region (f). To address these issues, we propose a novel boundary-limited shape optimization, which not only ensures the effectiveness of each shape, but also provides sufficient transformation space.

As shown in Figure 4 (b), the shaded part B_j is exactly the feasible region of the anchor point P_j . Specifically, the two adjacent equidistant lines, the inner circle's edge and the outer border together make up of our boundary. Figure 4 (c) presents us with an example of effective deformation. Next, we will give how to construct the boundary and the method of judging whether a point is inside or not.

We first divide a circle with a given radius r and a circle center C into n sharp-angled sectors by n equidistant lines $\{l_j | j = 1, \dots, n\}$. n equal points $\{E_j | j = 1, \dots, n\}$ are the intersection of equidistant lines and the circle, and as shown in Figure 4 (b), the anchor point P_j is the midpoint of E_j and E_{j+1} , which can be formulated as follows:

$$P_j = \frac{E_j + E_{j+1}}{2}, \quad j = 1, \dots, n \quad (8)$$

For the anchor point P_j , it lies in the sharp-angled sector region B_j wrapped by two adjacent lines l_j and l_{j+1} .




Figure 4. The solid blue and orange box represents the effective area of the human body and the dashed line means that the part is missing. The (a) expresses that the product of the directed distances from the point in the sharp-angled sector to the two side lines is less than 0 and if not, the product is more than 0. The (b) is our initial state and the (c) is an effective deformation. The (d) shows the patch contour will cross if we don't fix its relative order. The (e) represents that connecting points will be trapped into a restricted space if we don't set an inner circle. The (f) expresses that the patch contour will go beyond the human body if we don't define an outer boundary.

L_j and R_j represents the directed distance from the anchor point P_j to the line l_j and l_{j+1} shown in the Figure 4 (a). Therefore, with the line function $l_j(x, y) = 0$ and $l_{j+1}(x, y) = 0$, we can formulate L_j and R_j as follows:

$$L_j = \frac{l_j(x_j, y_j)}{D_j}, \quad R_j = \frac{l_{j+1}(x_j, y_j)}{D_{j+1}} \quad (9)$$

where D_j and D_{j+1} are the denominators of the point-to-line distance formula.

From the Figure 4 (b), we can vividly find that since the anchor point P_j 's feasible region is a sharp-angled sector, if P_j is inside its own region, $L_j * R_j < 0$ for the angle between them is obtuse and if P_j is outside its own region or exactly on the boundary line, $L_j * R_j \geq 0$. With P_j inside its own region, we can effectively avoid boundary line crossings in Figure 4 (d) caused by spline interpolation's need for a given order.

Then, to prevent points from falling into narrow places in Figure 4 (e), we set an inner circle with a given radius r^* inside the initial circle, and anchor points are not allowed to move into the inner circle. Specifically, we can judge by the distance r_j from the anchor point P_j to the center C . When $r_j > r^*$, the anchor point P_j will not be inside the inner circle. The r_j is computed as follows:

$$r_j = \sqrt{(x_j - x_C)^2 + (y_j - y_C)^2} \quad (10)$$

After that, we still need to ensure that the generated patch is not outside the effective area of the human body as shown in Figure 4 (f), so we scale the detection box output in a certain proportion as the outer border. The region \mathcal{O} inside the outer border can be represented as follows:

$$\mathcal{O} = \{(x, y) | x_l \leq x \leq x_r, y_d \leq y \leq y_u\} \quad (11)$$

where x_l, x_r, y_u, y_d are outer border's vertex coordinates.

Finally, we combine the above limits as follows:

$$\rho_j = (L_j * R_j < 0) \& (r_j > r^*) \& (P_j \in \mathcal{O}) \quad (12)$$

where $\rho_j = 1$ means that the anchor point P_j is inside the feasible region B_j and $\rho_j = 0$ means that the anchor point P_j is outside the boundary.

3.3.3 Score-aware Iterative Evaluation

In cross-modal attacks, a good attack effect of a single modality is ineffective, while it is common to have unbalanced attack effects in the two modalities. If this situation is not improved, it may give the attacker false signals about the progress of the attack. Thus, to balance the fooling degree between visible detector and infrared detector during the optimization process, we propose a score-aware iterative evaluation, guiding the adversarial patch to iteratively reduce predicted scores of the multi-modal sensors.

For the sake of simplicity of expression, here we denote $S_{ij}(k)$ as s . To evaluate the fitness value of s , we first use methods in Section 3.2 to transform s into a patch mask M , then $x_{vis}^{adv}, x_{inf}^{adv}$ are produced based on Eq.(2), Eq.(3). $J(s)$ can be formed as follows:

$$J(s) = e^{\lambda * \min(\text{dis}(x_{vis}^{adv}), \text{dis}(x_{inf}^{adv}))} \quad (13)$$

where λ is a weighted factor, $\text{dis}(x_{vis}^{adv}), \text{dis}(x_{inf}^{adv})$ reflect the current progress towards the success of attack (the larger, the closer to success). $\text{dis}(x_{vis}^{adv}), \text{dis}(x_{inf}^{adv})$ can be formalized as:

$$\text{dis}(x_{vis}^{adv}) = \frac{f_{vis}(x_{vis}) - f_{vis}(x_{vis}^{adv})}{f_{vis}(x_{vis}) - \text{thre}} \quad (14)$$

$$\text{dis}(x_{inf}^{adv}) = \frac{f_{inf}(x_{inf}) - f_{inf}(x_{inf}^{adv})}{f_{inf}(x_{inf}) - \text{thre}} \quad (15)$$

where $f_{vis}(x_{vis})$ is the confidence score of x_{vis} in the visible pedestrian detector, $f_{inf}(x_{inf})$ is the confidence score of x_{inf} in the infrared pedestrian detector. $f_{vis}(x_{vis}^{adv})$ and $f_{inf}(x_{inf}^{adv})$ are similar to $f_{vis}(x_{vis})$ and $f_{inf}(x_{inf})$.

From Eq.(14) and Eq.(15), we can easily know $\text{dis}(x_{vis}^{adv})$ and $\text{dis}(x_{inf}^{adv})$ measure the progress to success

of the cross-modal patch attack in the visible and infrared modality respectively. Both of them can help patches evolve in the corresponding modality. However, if we only make use of a single one, it will be certain to cause unbalanced phenomena, not an effective cross-modal attack. To solve this issue, we use Eq.(13) to combine $dis(x_{vis}^{adv})$ and $dis(x_{inf}^{adv})$. Based on our $J(\cdot)$, the good performance only in a single modality will not achieve a high value of fitness because we take the worse one of $dis(x_{vis}^{adv})$ and $dis(x_{inf}^{adv})$ as a standard. Additionally, considering the difference of attack difficulty in the initial stage and later stage, we use $e^{(\cdot)}$ instead of a linear function. Based on the above settings, we will eventually encourage it to iteratively evolve in the direction of reducing confidence scores as much as possible while maintaining the balance of cross-modalities.

The overall algorithm for generating cross-modal adversarial patches is summarized in Algorithm 1.

Algorithm 1: Shape-based Cross-modal Attack

```

Input: Clean visible image  $x_{vis}$ , clean infrared
          image  $x_{inf}$ , fitness function  $J(\cdot)$ , population
          size  $Q$ , max iteration number  $T$ 
1 Initialize a collection of shapes  $S(0)$ 
  /* Optimize patch shape with DE          */
2 for  $k = 0$  to  $T - 1$  do
3   Sort  $S(k)$  in descending order based on  $J(S(k))$ 
4   if  $S_0(k)$  makes the attack successful then
5     |  $stop = k$ ; break;
6   end
7   Generate  $S(k + 1)$  based on crossover and mutation
8   Limit boundaries of  $S(k + 1)$  according to Eq.(12)
9   for  $i = 1$  to  $Q$  do
10    | Evaluate  $S_i(k)$ ,  $S_i(k + 1)$  according to Eq.(13)
11    |  $S_i(k + 1) \leftarrow$  the better one in  $S_i(k)$ ,  $S_i(k + 1)$ 
12    end
13 end
14 Sort  $S(stop)$  in descending order according to  $J(\cdot)$ 
  /* Model patch shape from points          */
15 Anchor points  $\{P_{i=1, \dots, n}\} \leftarrow S_0(stop)$ 
16 Curve segments  $\{C_{i=1, \dots, n}\} \leftarrow$  connect  $\{P_{i=1, \dots, n}\}$ 
   according to Eq.(4)
17  $M_{con} \leftarrow$  combine  $\{C_{i=1, \dots, n}\}$  according to Eq.(5)
18  $M \leftarrow$  fill the  $M_{con}$  according to Eq.(6)
19  $x_{vis}^{adv}$ ,  $x_{inf}^{adv} \leftarrow$  according to Eq.(2), Eq.(3)
20 return  $M$ ,  $x_{vis}^{adv}$ ,  $x_{inf}^{adv}$ 
Output: Mask  $M$ , adversarial visible example  $x_{vis}^{adv}$  and
          adversarial infrared example  $x_{inf}^{adv}$ 

```

3.4. Physical Implementation

After obtaining the optimal shape, we start to transform algorithm-generated digital patches into physical patches with the aerogel material. Specifically, we first scale the obtained patch according to its size in the real world and

print it. Then, we use scissors to equally cut out the optimal shape on the material. And in the last, we use the velcro sticker to fix it in the corresponding position of the human body for performing a physical attack. The whole process is demonstrated in Figure 5.



Figure 5. The physical implementation process of transforming digital cross-modal patches into physical patches in the real world.

4. Experiments

4.1. Simulation of Physical Attacks

Dataset: We use the LLVIP [9] dataset to simulate the physical attacks. Images from LLVIP are perfectly synchronized in both visible modality and infrared modality. Similar to [34] and [33], we customize the parts containing pedestrians from images in LLVIP. There are 1220 photos in the test set and 3784 images in the training set. As the final samples to be attacked, we pick 120 images from the test set that the target model can recognize successfully with a high probability. The clean AP is therefore 100%.

Target detector: For the pedestrian detection task, we select two mainstream detectors: YOLOv3 [14](one-stage) and Faster RCNN [15](two-stage) here. About model training, we choose the officially pre-trained weights as the initialized weights and then retrain the model on the training dataset. Results of some other detectors like YOLOv5, YOLOv7 [23], SSD [12], and EfficientDet[19] are shown in the *Supplementary Material*.

Metrics: Attack Success Rate (ASR) and Average Precision drop (AP drop) are used to evaluate the attack performance. Here, we adopt a unique cross-modal ASR, which denotes the ratio of both successfully attacked images under two modalities out of all the test images, highlighting the effectiveness of simultaneously attacking two modalities. AP drop is to show the AP's variation before and after attacks.

Implementation: As we use the DE algorithm, we set the number of the initial population as 30, the epoches of evolution as 200. Other parameters, formulas and comparative experiments are presented in the *Supplementary Material*.

4.1.1 Performances in Different Detection Systems

We first evaluate the effectiveness of attacks in the digital world. Considering the differences between detection systems, we choose two typical detectors: YOLOv3 and Faster

Table 2. Attack performances in different detection systems.

	YOLOv3	Faster RCNN
ASR	73.33%	69.17%
AP drop (Visible)	99.19%	89.38%
AP drop (Infrared)	74.31%	83.94%

RCNN. The results are shown in Table 2. We use ASR and AP drop to evaluate the attack performance.

From the above results, we can see that our method is equally useful despite the distinctions between the one-stage and two-stage detection models. For the typical one-stage detection model YOLOv3, we achieve an ASR of 73.33%, an AP drop of 99.19% in the visible modality and 74.31% in the infrared modality. Similarly, for the two-stage detection model Faster RCNN, we achieve an ASR of 69.17%, an AP drop of 89.38% in the visible modality and 83.94% in the infrared modality.



Figure 6. Visual examples of adversarial samples with cross-modal patches in the digital world.

4.1.2 Effects of Optimized Shapes

Here, we provide the ablation study to investigate the outcomes of our optimized shapes. Specifically, we first generate cross-modal adversarial patches to achieve the optimal shapes for a given pedestrian. Some examples of the optimal shapes in the digital world are shown in Figure 6. Then, we fix the patches' locations and sizes while changing their shapes. These shapes can be circles, squares, rectangles and triangles, etc. Then we compute the ASR and AP drop under these five shapes. This setting tests the effects of optimal shapes on the attacks. The results are listed in Table 3, where we can see that random shapes barely work with an average ASR of 5.23%, an average AP drop of 20.73% in the visible modality and 11.65% in the infrared modality. The results for each of the five shapes and compared with [25] are shown in the *Supplementary Material*.

Table 3. Ablation study for unified adversarial patches' shapes.

	Our shape	Other shapes
ASR	73.33%	5.23%
AP drop (Visible)	99.19%	20.73%
AP drop (Infrared)	74.31%	11.65%

4.1.3 Effects of Score-aware Iterative Fitness Function

To verify the impact of the Cross-modal Fitness Function, we design a set of comparative experiments. As mentioned in Section 3.3.3, our method can combine the visible modality and infrared modality into a whole, which helps balance attack performances between different modalities. Here, we use a simple sum of Eq.(14) and Eq.(15) instead of Eq.(13) as the not-combined fitness function.

Table 4. Ablation study for score-aware iterative fitness function.

	Ours	Sum
ASR	73.33%	50.83%
AP drop (Visible)	99.19%	85.37%
AP drop (Infrared)	74.31%	66.99%

From Table 4, we can find that although a simple sum fitness function can have an AP decrease of 85.37% in the visible modality and 66.99% in the infrared modality, it suffers a decline in ASR under a cross-modal standard.

4.1.4 Hyperparameters Tuning

Patch number. As one of the most directly adjustable hyperparameters, patch number may have an impact on the attack performance. Therefore, we need to explore whether and how it is related to the attack performances. Figure 7 (a) vividly shows that as the number of patches rises, both ASR and AP drop increase with our expectation. However, considering that our attack method has achieved acceptable results with two patches, and three patches would bring larger physical implementation errors, we finally decide to choose the patch number 2 as the main evaluation option.

Hyperparameter λ . The λ in Eq.(13) affects the variability of the final fitness of the individual, the larger the λ , the greater the difference between different individuals, but this does not mean that the greater the difference, the better the attack effects. As shown in Figure 7 (b), the attack effectiveness of cross-modal patches increases when λ equals from 1 to 2, but drops dramatically when λ equals 3. As a result, we set $\lambda = 2$ to obtain the optimal effect.

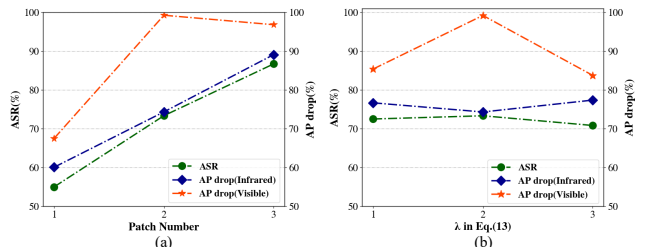


Figure 7. The ASR(%) and AP drop(%) of unified adversarial patches with different patch numbers and hyperparameter λ .

4.1.5 Robustness to Implementation Errors

Since the patch is in general generated digitally by considering the optimal shape and fixed location, when it is applied

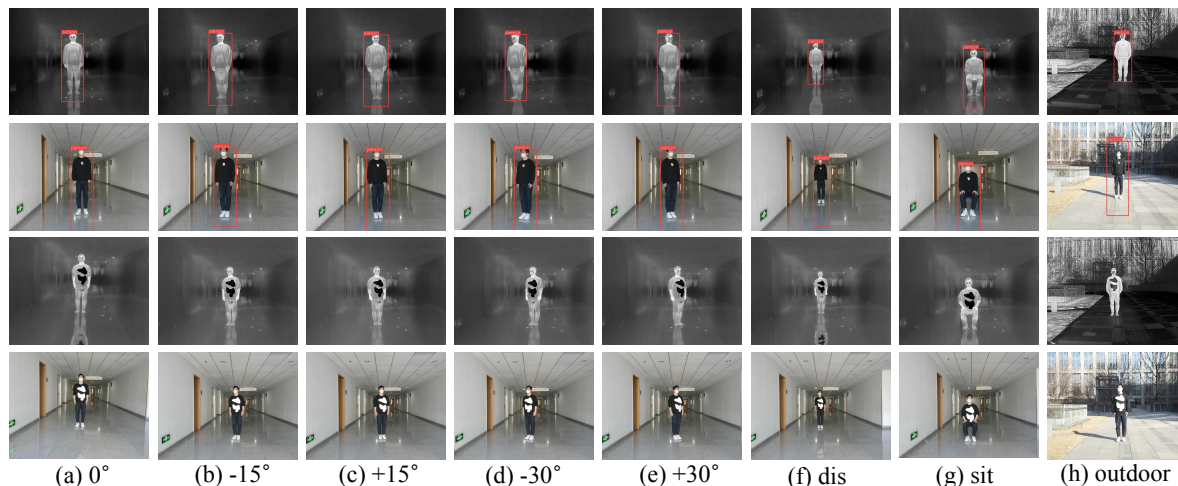


Figure 8. Visual examples of physical attacks with cross-modal patches under various angles, distances, postures, and scenes.

to real scenarios, it is natural to see how sensitive it would be to the attack success rate if we do not cut and paste the patch in 100% exact shape and location on the clothes.

Therefore, we simulate the possible position shifts and clipping errors when doing physical experiments. Table 5 demonstrates our cross-modal patches' robustness to translation errors and incompleteness.

Table 5. ASR of simulating implementation errors.

	Translation		Incompleteness	
	3pix	5pix	5%	10%
ASR	64.17%	55.00%	68.83%	60.83%
	(↓9.16%)	(↓18.33%)	(↓5.00%)	(↓2.50%)

4.2. Attacks in the physical world

To test the effectiveness of our method in the real world, we build a series of scenarios and record actual videos under corresponding physical settings to compute their ASR. By default, the recordings of interior scenes are conducted 4 meters away from a standing person in frontal perspective (0°), and we record each scenario for 20 seconds at a rate of 10 frames per second (about 200 frames in total). Similar to [33], we set the threshold of pedestrian detection as 0.7. For different settings, in the angle issue, we use $\pm 15^\circ$ and $\pm 30^\circ$ for verification. In the distance issue, we change the camera's default setting from 4 meters to 6 meters. In the posture issue, the pedestrian's posture is replaced by a sitting one. In the scene issue, we switch from the default indoor to the outdoor. Figure 8 lists the concrete illustrations of various scenarios. Besides, the pedestrian is asked to move the body within the range of 5° of the current posture to take videos while shooting. Quantitative results are shown in Table 6. It is clear that from the frontal perspective, our unified adversarial patches could obtain a high ASR (73.50%). The ASR still maintains a high value (67.00% and 46.50%) when the shooting angle varies. The ASR drops to 71.00%

when the distance is increased to 6 meters from 4 meters. The ASR drops to 61.50% when the posture is changed from standing to sitting, and to 57.00% at exterior scenes. These findings demonstrate that our patches are not significantly impacted by various shooting scenarios. In other words, the effects of adversarial attacks can be maintained as long as the camera can capture the whole shape of our cross-modal patches on the object.

Table 6. ASR in the physical world when changing angles, distances, postures and scenes captured by multi-modal sensors.

	0°	$\pm 15^\circ$	$\pm 30^\circ$	dist.	pos.	outdoor
ASR	73.50%	67.00%	46.50%	71.00%	61.50%	57.00%

4.3. Defenses against Unified Adversarial Patches

To verify the robustness towards adversarial defense, we utilize two common methods in the digital world: spatial smoothing [32] as a pre-processing defense and adversarial training [8] as a post-processing defense. The original ASR is 73.33% and the results in Table 7 shows that: (1) ASR only decreases by 9.16% after spatial smoothing. This is reasonable since our cover image have the same value that could not be altered by smoothing. (2) In the second case, ASR drops by 23.33%, which is still within an acceptable range. The whole results prove our method's robustness.

Table 7. Results against the defense methods.

Defense Methods	ASR	Error
Spatial Smoothing [32]	64.17%	9.16%
Adversarial Training [8]	50.00%	23.33%

5. Conclusion

In this paper, we propose a unified adversarial patch in the physical world. For that, we uncover the property that can react both in the visible and infrared modalities: **shape**.

Then, combining the boundary-limited shape optimization with the score-aware iterative fitness evaluation, we guarantee an efficient exploration of the adversarial shape and the balance between different modalities. Experiments on the pedestrian detection in the digital world and physical world verify the effectiveness of our method.

Acknowledgment

This work was supported by the Project of the National Natural Science Foundation of China (No.62076018), and the Fundamental Research Funds for the Central Universities.

References

- [1] Mazen Abdelfattah, Kaiwen Yuan, Z Jane Wang, and Rabab Ward. Adversarial attacks on camera-lidar models for 3d car detection. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2189–2194. IEEE, 2021. **1**
- [2] Tom B. Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. *CoRR*, abs/1712.09665, 2017. **2**
- [3] Tom B Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. *arXiv preprint arXiv:1712.09665*, 2017. **2**
- [4] 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. **1**
- [5] Zhaoyu Chen, Bo Li, Shuang Wu, Jianghe Xu, Shouhong Ding, and Wenqiang Zhang. Shape matters: deformable patch attack. In *European Conference on Computer Vision*, pages 529–548. Springer, 2022. **2, 3**
- [6] Yinpeng Dong, Shouwei Ruan, Hang Su, Caixin Kang, Xingxing Wei, and Jun Zhu. Viewfool: Evaluating the robustness of visual recognition to adversarial viewpoints. *Advances in Neural Information Processing Systems*, 35:36789–36803, 2022. **1**
- [7] 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. **2**
- [8] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. **1, 8**
- [9] Xinyu Jia, Chuang Zhu, Minzhen Li, Wenqi Tang, and Wenli Zhou. Lvip: A visible-infrared paired dataset for low-light vision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3496–3504, 2021. **6**
- [10] Stepan Komkov and Aleksandr Petiushko. Advhat: Real-world adversarial attack on arcfac face id system. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 819–826. IEEE, 2021. **1**
- [11] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016. **1, 2**
- [12] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 21–37. Springer, 2016. **6**
- [13] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. **1**
- [14] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. **6**
- [15] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. **6**
- [16] 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. **2**
- [17] Jiachen Sun, Yulong Cao, Qi Alfred Chen, and Z Morley Mao. Towards robust {LiDAR-based} perception in autonomous driving: General black-box adversarial sensor attack and countermeasures. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 877–894, 2020. **1**
- [18] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. **1**
- [19] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficientdet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10781–10790, 2020. **6**
- [20] Simen Thys, Wiebe Van Ranst, and Toon Goedemé. Fooling automated surveillance cameras: adversarial patches to attack person detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019. **1**
- [21] James Tu, Huichen Li, Xinchun Yan, Mengye Ren, Yun Chen, Ming Liang, Eilyan Bitar, Ersin Yumer, and Raquel Urtasun. Exploring adversarial robustness of multi-sensor perception systems in self driving. *arXiv preprint arXiv:2101.06784*, 2021. **1**
- [22] Christopher Twigg. Catmull-rom splines. *Computer*, 41(6):4–6, 2003. **4**
- [23] Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors, 2022. **6**
- [24] Shaojie Wang, Tong Wu, and Yevgeniy Vorobeychik. Towards robust sensor fusion in visual perception. *CoRR*, abs/2006.13192, 2020. **1**
- [25] Hui Wei, Zhixiang Wang, Xuemei Jia, Yinqiang Zheng, Hao Tang, Shin’ichi Satoh, and Zheng Wang. Hotcold block:

- Fooling thermal infrared detectors with a novel wearable design. *arXiv preprint arXiv:2212.05709*, 2022. [1](#), [2](#), [7](#)
- [26] 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. [1](#)
- [27] 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. [1](#)
- [28] Xingxing Wei, Huanqian Yan, and Bo Li. Sparse black-box video attack with reinforcement learning. *International Journal of Computer Vision*, 130(6):1459–1473, 2022. [1](#)
- [29] Xingxing Wei, Jie Yu, and Yao Huang. Physically adversarial infrared patches with learnable shapes and locations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12334–12342, June 2023. [1](#)
- [30] Xingxing Wei, Jun Zhu, Sha Yuan, and Hang Su. Sparse adversarial perturbations for videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8973–8980, 2019. [1](#)
- [31] Kaidi Xu, Gaoyuan Zhang, Sijia Liu, Quanfu Fan, Mengshu Sun, Hongge Chen, Pin-Yu Chen, Yanzhi Wang, and Xue Lin. Adversarial t-shirt! evading person detectors in a physical world. In *European conference on computer vision*, pages 665–681. Springer, 2020. [1](#), [2](#)
- [32] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155*, 2017. [8](#)
- [33] Xiaopei Zhu, Zhanhao Hu, Siyuan Huang, Jianmin Li, and Xiaolin Hu. Infrared invisible clothing: Hiding from infrared detectors at multiple angles in real world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13317–13326, 2022. [1](#), [2](#), [6](#), [8](#)
- [34] 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. [1](#), [2](#), [6](#)