

PBCAT: Patch-Based Composite Adversarial Training against Physically Realizable Attacks on Object Detection

Xiao Li^{1*} Yiming Zhu^{1,2*} Yifan Huang^{1,2*} Wei Zhang¹ Yingzhe He³ Jie Shi³ Xiaolin Hu^{1,4†}

¹ Department of Computer Science and Technology, BNRist,
IDG/McGovern Institute for Brain Research, THBI, Tsinghua University

²University of Science and Technology Beijing

³Huawei Technologies ⁴Chinese Institute for Brain Research (CIBR)

{lixiao20, zhang-w19}@mails.tsinghua.edu.cn, {U202241171, U202240988}@xs.ustb.edu.cn,
{heyinzhe, shi.jie1}@huawei.com, xlhu@mail.tsinghua.edu.cn

Abstract

Object detection plays a crucial role in many security-sensitive applications. However, several recent studies have shown that object detectors can be easily fooled by physically realizable attacks, e.g., adversarial patches and recent adversarial textures, which pose realistic and urgent threats. Adversarial Training (AT) has been recognized as the most effective defense against adversarial attacks. While AT has been extensively studied in the l_∞ attack settings on classification models, AT against physically realizable attacks on object detectors has received limited exploration. Early attempts are only performed to defend against adversarial patches, leaving AT against a wider range of physically realizable attacks under-explored. In this work, we consider defending against various physically realizable attacks with a unified AT method. We propose PBCAT, a novel Patch-Based Composite Adversarial Training strategy. PBCAT optimizes the model by incorporating the combination of small-area gradient-guided adversarial patches and imperceptible global adversarial perturbations covering the entire image. With these designs, PBCAT has the potential to defend against not only adversarial patches but also unseen physically realizable attacks such as adversarial textures. Extensive experiments in multiple settings demonstrated that PBCAT significantly improved robustness against various physically realizable attacks over state-of-the-art defense methods. Notably, it improved the detection accuracy by 29.7% over previous defense methods under one recent adversarial texture attack¹.

*Equal contribution.

†Corresponding author.

¹Code is available at <https://github.com/LixiaoTHU/oddefense-PatchAT>

1. Introduction

Object detection, which requires simultaneously classifying and localizing all objects in an image, is a fundamental task in computer vision. Recent advancements in object detection methods [25, 39, 46, 59] have greatly benefited from the utilization of Deep Neural Networks (DNNs). However, DNNs are known to be susceptible to adversarial examples [43] crafted by adding deliberately designed perturbations to the original examples.

Adversarial examples exist not only in the digital world but also in the physical world [3, 16, 17, 45, 51, 52]. Several studies have demonstrated that object detectors can be easily fooled by *physically realizable attacks*, e.g., adversarial patches [3, 45] and adversarial textures [16, 17]². Specifically, adversarial *patch attacks* craft localized adversarial patterns within a fixed region (e.g., a square patch), while adversarial *texture attacks* craft more pervasive adversarial perturbations that spread across the entire surface of the object, e.g., adversarial modifications to clothing textures that cover most of the surface of an object. Given the crucial role of object detection in numerous security-sensitive real-world applications, such as autonomous driving [1] and video surveillance [22], it is imperative to improve the adversarial robustness of object detectors against these physically realizable attacks, which pose realistic and severe threats.

Numerous studies [21, 32, 34–36, 50, 54, 55, 60] have proposed various defense methods against physically realizable attacks. However, three critical limitations still hinder their effectiveness: 1) Most existing defenses are designed specifically for patch attacks, which represent only the simplest form of physically realizable attacks. Recent

²Both adversarial patches and adversarial textures can be implemented in the physical world and are thereby called physically realizable attacks.

adversarial texture attacks [16, 17], which generate adversarial clothing to deceive person detectors in the physical world, introduce a distinct threat model that differs from adversarial patches. To the best of our knowledge, defenses against such attacks remain largely unexplored. 2) Existing methods may be circumvented by specially designed techniques, *i.e.*, adaptive attacks [2, 47]. Adversarial Training (AT) [33] is widely recognized as an effective defense against adaptive attacks and has been extensively studied in the l_∞ attack setting (referred to as l_∞ -AT) [6, 26, 28–30, 33, 56, 58]. However, l_∞ attacks involve global adversarial perturbations applied to all image pixels, which are impractical in the physical world. Thus, l_∞ -AT, which is designed for a fundamentally different threat model, may perform poorly against physically realizable attacks. 3) AT with adversarial patches (termed *patch-based AT*) is a promising approach for defending against patch attacks. However, only a few early works [34, 36, 50] have explored its application in classification tasks. Extending these methods to object detection incurs prohibitive computational costs, which we discuss in Sec. 3.3.

To address these limitations, we consider defending against various physically realizable attacks with a unified AT method. Specifically, we introduce PBCAT, a novel Patch-Based Composite Adversarial Training strategy. First, we extend l_∞ -AT to patch-based AT, to improve robustness against adversarial patch attacks. Second, we propose a gradient-guided patch partition and selection method to efficiently find effective patch locations for AT. Third, we incorporate global imperceptible adversarial perturbations used in l_∞ -bounded AT into patch-based AT, which primarily employs small-area, gradient-guided adversarial patches. By leveraging composite perturbations, PBCAT mitigates overfitting to patch attacks, enabling it to defend not only against adversarial patches but also against unseen physically realizable attacks, such as adversarial texture attacks with large-area perturbations. Finally, to enhance the practical utility of PBCAT, we draw inspiration from FreeAT [41] to enable PBCAT to train a robust detector at a computational cost comparable to standard training.

We trained object detectors with PBCAT on the MS-COCO [31] dataset. Evaluations were performed on several datasets, including the MS-COCO dataset for the general object detection task as well as the Inria [10] dataset along with a synthetic dataset [16] for the downstream security-critical person detection task. We demonstrated that PBCAT significantly improved robustness against various physically realizable attacks over state-of-the-art (SOTA) defense methods in strong adaptive settings [2, 47]. On the person detection task, PBCAT secured a Faster R-CNN [39] with 60.2% and 56.4% Average Precision (AP) against two recent adversarial texture attacks, AdvTexture [16] and AdvCaT [17], respectively. Notably, PBCAT achieved a 29.7%

AP improvement over the SOTA defense methods against AdvTexture.

Our main contributions can be summarized as follows:

- We propose PBCAT, a novel adversarial training method that unifies the defense against various physically realizable attacks within a single model;
- PBCAT closes the gap between adversarial patches and adversarial textures by employing composite perturbations along with patch partition and gradient-guided selection techniques;
- Extensive experiments demonstrate that PBCAT achieves superior adversarial robustness against diverse physically realizable attacks in strong adaptive settings.

2. Preliminary and Related Work

2.1. Adversarial Robustness on Classification

Adversarial examples, first discovered on image classification [43], are input images with deliberately designed perturbations that can fool DNN-based image classifiers while still being easily recognized by humans. Given an image-label pair (\mathbf{x}, y) and a classifier $f_\theta(\cdot)$, adversarial perturbation δ can be found by maximizing the output loss: $\delta = \arg \max_{\mathcal{B}(\delta) \leq \epsilon} \mathcal{L}(f_\theta(\mathbf{x} + \delta), y)$, where \mathcal{L} denotes a cross-entropy loss, and the attack intensity ϵ bounds the attack budget B . Several approximate methods [4, 13, 33] have been proposed to solve the intractable maximizing problem. PGD [33] is one of the most popular methods, which optimizes perturbations through multiple iterations with small step sizes. AT and its variants are generally considered the most effective defense methods against adversarial examples, which improve adversarial robustness by incorporating adversarial examples into training:

$$\theta = \arg \min_{\theta} \mathbb{E}_{\mathbf{x}} \left\{ \max_{\mathcal{B}(\delta) \leq \epsilon} \mathcal{L}(f_\theta(\mathbf{x} + \delta), y) \right\}. \quad (1)$$

However, most works investigate AT on classification models in the l_∞ -bounded settings [28, 33, 58], *i.e.*, $\mathcal{B}(\cdot) := \|\cdot\|_\infty$, which involve adding a global adversarial perturbation to the images and are generally considered as physically infeasible attacks.

2.2. Adversarial Attacks on Object Detection

Object detection requires simultaneously classifying and localizing all objects in an image. Modern object detectors [39, 46, 59] have significantly improved with the utilization of DNNs. Object detectors are also vulnerable to adversarial examples, and several works [6, 9, 12] have investigated how to attack object detectors from various aspects. Unlike classifiers where l_∞ attacks and defenses are often investigated, more urgent physically realizable attacks are widely studied for this task, considering that object detection has been widely used in many security-critical ap-

plications. Particularly, for the person detection task, several physically realizable adversarial attacks have been proposed. Thys et al. [45] first propose to generate physically realizable adversarial patches to fool person detectors, which we denote AdvPatch. AdvTexture attack [16] further extends AdvPatch to tileable adversarial textures, offering attack effects from various viewing angles. It proposes a scalable generative method to craft adversarial texture with repetitive structures. AdvCaT attack [17] optimizes adversarial textures into typical camouflage patterns to resemble cloth patterns in the physical world. Attackers can print the clothing texture created by AdvCaT and AdvTexture on a piece of cloth and tailor it into an outfit. Wearing such an outfit can hide the person from SOTA detectors, posing realistic and urgent security threats.

2.3. Adversarial Defenses against Patch Attacks

To defend against adversarial examples of object detectors, especially patch-based attacks, several types of methods have been proposed. Input preprocessing-based methods, such as LGS [35], SAC [32], EPGF [60], and Jedi [44], mask out or suppress the potential adversarial patch areas before sending them to the model. Outlier feature filter-based methods, such as FNC [54] and APE [21], incorporate filters into the model to smooth abnormal inner features caused by adversarial patches. Defensive frame methods, such as UDF [55], train an adversarial defense frame surrounding images to improve robustness. However, similar to the experiences on l_∞ defense [2, 47], These non-AT methods can be vulnerable to strong adaptive attacks, as detailed in our experimental results in Sec. 4.2. Besides these empirical methods, some studies [7, 8] investigate to improve certified robustness, but till now the certified methods only work under quite tiny perturbations and need time-consuming inference. In this work, we mainly compare PBCAT with empirical methods in practical scenarios.

Only a few early works [34, 36, 50] have explored patch-based AT for defending against patch attacks. Both Rao et al. [36] and Wu et al. [50] investigate methods for identifying optimal patch locations for patch-based AT in classification models, while Metzen et al. [34] propose an enhanced patch-based AT approach based on meta-learning. Although these methods demonstrate promising results against patch attacks, they were not originally designed for object detection, and adapting them to this task presents significant challenges. For instance, Rao et al. [36] and Wu et al. [50] primarily focus on selecting a single optimal patch position for classification, which typically involves a single object. In contrast, object detection requires handling multiple bounding boxes, necessitating the identification of multiple effective patch locations. As a result, directly applying these methods to object detection leads to a substantial increase in computational complexity. More-

over, relying solely on adversarial patches for training does not generalize well to a broader range of physically realizable attacks, as discussed later.

3. PBCAT

PBCAT aims to defend against various physically realizable attacks by incorporating a combination of a small area of sub-patches and a large area of imperceptible adversarial perturbations. We discuss how to build patch-based AT methods from l_∞ -AT in Sec. 3.1. We then describe the gradient-guided patch partition and selection method in Sec. 3.3. In Sec. 3.4, we introduce how to extend patch-based AT to defend against large-area perturbations. Finally, we show how to accelerate PBCAT in Sec. 3.5.

3.1. From l_∞ -AT to Patch-Based AT

Given an image $\mathbf{x} \in [0, 1]^{3 \times H \times W}$ and its corresponding bounding box labels y , where $H \times W$ denotes the input resolution, l_∞ attacks typically constrain the perturbation budget \mathcal{B} with a bound ϵ : $\|\cdot\|_\infty \leq \epsilon$. In practice, ϵ is often set to a small value, e.g., $4/255$ [29, 40]. However, l_∞ -bounded perturbations apply *global* modifications to all pixels in \mathbf{x} , which is generally considered physically unrealistic. In contrast, physically realizable attacks restrict perturbations to *local* regions, typically targeting the foreground of an object [16, 17, 45], where pixel modifications can have a higher intensity. Unlike classification tasks, object detection involves multiple foreground objects within a single input image. In this work, we adopt a threat model where each bounding box may contain an adversarial patch. Since l_∞ attacks and physically realizable attacks assume different perturbation constraints, and AT often exhibits suboptimal performance against unseen threats [24, 57], we opt for patch-based AT as a more effective defense against physically realizable attacks.

On the other hand, according to Eq. (1), patch-based AT follows a similar formulation to l_∞ -AT, with the key distinction being the introduction of a mask to restrict perturbations to specific regions. Specifically, patch-based AT for object detectors can be formulated as follows:

$$\theta = \arg \min_{\theta} \mathbb{E}_{\mathbf{x}} \left\{ \max_{\|\delta_p \odot \mathbf{M}\|_\infty \leq \beta} \mathcal{L}_d(f_\theta(\mathbf{x} + \delta_p \odot \mathbf{M}), y) \right\}, \quad (2)$$

where δ_p denotes the patch perturbation, while \mathbf{M} denotes a binary mask for restricting δ_p to a local area. β denotes a large perturbation intensity (e.g., 1), and \mathcal{L}_d denotes the loss of an object detector. The inner maximizing problem of patch-based AT can take inspiration from l_∞ -AT, as detailed in Sec. 3.5. We discuss how to obtain \mathbf{M} next.

3.2. Instance-Level Patch Placement

For object detection tasks, the adversarial patch is typically created and placed to a fixed location relative to the bound-

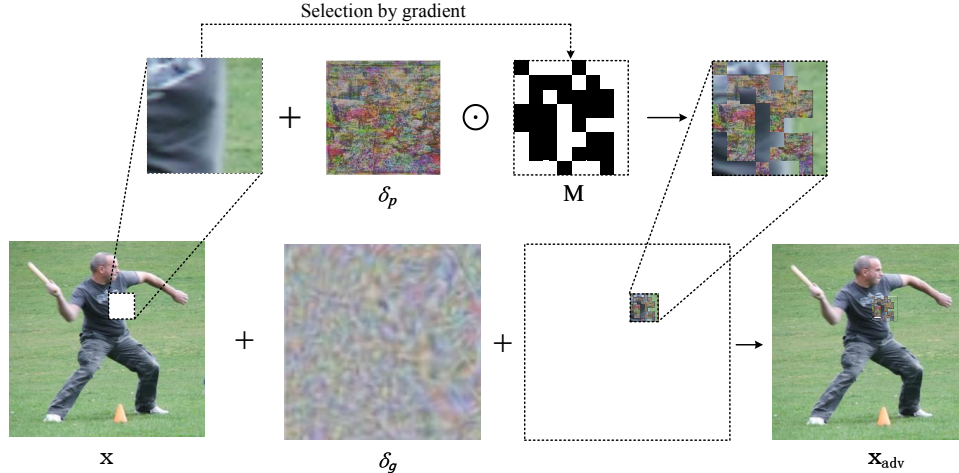


Figure 1. The illustration on how to generate training images for PBCAT. A patch location is randomly selected from each bounding box area first. The selected patch is then partitioned into multiple sub-patches. The average gradient norm is computed for each partitioned sub-patch, and the top half with the highest values are selected to construct a binary mask M . This mask is then applied element-wise to the patch perturbation δ_p , forming a new masked patch. The global small perturbation δ_g and these sub-patches are added to obtain the final adversarial example x_{adv} . Since δ_g is constrained by $\|\delta_g\|_\infty \leq 4/255$, we scale δ_g in the visualization for better clarity.

ing box. But adopting this approach in AT may cause information leakage for bounding box prediction, as models might utilize the adversarial patch to perform bounding box regression. PBCAT randomly samples adversarial patch locations within a bounding box to mitigate this issue. Denoting the size of the bounding box as (w_{bbox}, h_{bbox}) , and the center of the bounding box as (x_{bbox}, y_{bbox}) , the center (x_s, y_s) of the patch is obtained by randomly sampling patch locations within the bounding box via a Gaussian distribution with (x_{bbox}, y_{bbox}) as the mean. The width and height of the sampled patch are set to be $s = \lambda \cdot \sqrt{w_{bbox}^2 + h_{bbox}^2}$, where λ is a hyper-parameter.

3.3. Gradient-Guided Patch Partition and Selection

As highlighted in early patch-based AT attempts [36, 50], patch location plays a crucial role in enhancing the effectiveness of patch-based AT. However, these methods often rely on time-consuming trial-and-error processes, such as multiple forward inferences to determine optimal patch locations [36] within a specific optimization step. This issue is exacerbated in object detection, where an image typically contains multiple objects, leading to a computational overhead that scales with the number of candidate locations. Thus, extending these methods to object detection incurs prohibitive theoretical computational costs, which we further analyze in Appendix A.

In contrast, PBCAT efficiently determines patch locations using gradient-based selection, requiring only a single forward/backward pass for each optimization step. For simplicity, here we consider the case of a single bounding box, noting that the multi-object scenario follows a similar process. Given a randomly placed patch (as described in

Sec. 3.2), we partition it into $n \times n$ sub-patches, resulting in $N = n^2$ partitioned areas, where N is a hyper-parameter. After a single backward pass to obtain the whole gradient information for this patch, we compute the average gradient norm for each sub-patch and select the top 50% with the highest values, as the areas with large gradient norms generally are the vulnerable areas that have a significant impact on the output loss. For implementation, we define a binary mask M of the same size as the original patch. This mask, applied via element-wise multiplication, determines which regions of the patch are retained. An illustration of M is shown at the top of Fig. 1. Our method effectively identifies vulnerable regions for adversarial training using only a single forward and backward pass, incurring negligible additional computational cost compared to previous methods.

3.4. Local Patches and Global Noises

Recent adversarial texture attacks [16, 17] adopt a significantly larger area attack than patch-based attack [45]. Training with small adversarial patches only makes it challenging to defend against these large-area physically realizable attacks (see Sec. 4.3, where training with patches alone exhibited poor robustness against texture attacks). To defend against these attacks, a direct and ideal method would be to perform AT with large-area unrestricted adversarial noises. However, we find that simply increasing the patch size (perturbation area) can induce training collapse of patch-based AT, *i.e.*, the large-area unrestricted perturbation incurred slow training convergence and poor robustness (see Tab. 5, where doubling the patch size significantly reduced robustness against various attacks). We guess that the collapse might be caused by the model’s limited capacity and sig-



Figure 2. Visualization of training examples with adversarial perturbations used in PBCAT. Perceptible gradient-guided adversarial patches and imperceptible global adversarial perturbations are added to each image. The patch regions are annotated with white dashed boxes.

nificant corruption of object information (a large-area unrestricted patch perturbation can corrupt the information of the entire object).

Instead, we propose to incorporate global imperceptible adversarial perturbations generally used in l_∞ -AT into the patch-based AT. The insight behind this approach is as follows: 1) By incorporating large-area adversarial noises while constraining the attack intensity (l_∞ -bound), sufficient object information can be kept to avoid training collapse. 2) On the other hand, since the perturbation areas of unseen physically realizable attacks may be located at arbitrary locations of an image (e.g., using a printable patch), the l_∞ -bounded global noise ensures the entire image to be covered by adversarial perturbations during training. 3) Training with l_∞ -bounded noise has been shown to be helpful against several different adversarial threats [27, 48], and our subsequent results further confirms that l_∞ -AT can enhance robustness against patch attacks. The final perturbation used in PBCAT is:

$$\delta = \text{Apply}(\delta_p \odot \mathbf{M}, \mathbf{x}) + \delta_g, \quad (3)$$

where δ_p denotes the patch perturbation, δ_g denotes the global noises, i.e., $\|\delta_g\|_\infty \leq \epsilon$, and Apply denotes the operation of pasting the patch back onto the image \mathbf{x} .

3.5. Accelerating AT

Early patch-based AT works [34, 36, 50] employ full PGD attacks to solve the inner maximization problem of AT and train object detectors from scratch, making the process highly time-consuming. In PBCAT, inspired by recent SOTA l_∞ -AT practices [29] for object detection, we adopt FreeAT [41] as the default setting for patch-based AT and utilize an adversarially pre-trained backbone network [29]. FreeAT recycles gradient perturbations to mitigate the additional training cost introduced by inner maximization while maintaining comparable adversarial robustness. We initialize both the patch perturbation δ_p and the global perturbation δ_g to zero. At each iteration, we compute the gradient, take its sign, and update δ_p and δ_g using different step sizes. The pseudo-code of “Free” PBCAT for object detection is provided in Appendix B. Notably, the adversarial gradient

\mathbf{g}_{adv} and model parameter gradient \mathbf{g}_θ are computed simultaneously in a single backward pass. By leveraging “Free” PBCAT, the training cost of adversarial training becomes comparable to standard training. The actual training time is detailed in Appendix C.

4. Experiments

4.1. Experimental Settings

Unless otherwise specified, our primary experiments were conducted on the widely used two-stage detector Faster R-CNN [39] with a ResNet-50 [14] backbone. Additionally, we evaluated PBCAT on the one-stage detector FCOS [46] and the recent DN-DETR [25], with results presented in Sec. 4.4. Notably, Faster R-CNN, FCOS, and DN-DETR have distinct architectures. We did not include YOLO [20, 38] series detectors in this study, as discussed in Appendix C. PBCAT was trained on a diverse set of general objects. However, since most physically realizable attacks [3, 16, 17, 45] target security-critical applications, such as person detection, we mainly evaluated PBCAT-trained detectors on the person detection task.

Datasets and metrics. Three datasets are used in this work: 1) The MS-COCO [31] dataset for training the general object detectors. We used the 2017 version, containing 118,287 images of 80 object categories for training and 5,000 images for evaluation. 2) The Inria Person [10] dataset for the person detection task. Note that this dataset is only used for attack evaluation. 3) The synthetic dataset used in Hu et al. [16, 17] for evaluation. This dataset contains 506 background images of different scenes, with 376 images used for adversarial texture optimization and 130 images for evaluation. We rendered a 3D person wearing an adversarial outfit optimized by AdvTexture and AdvCaT on the provided background images using a differential renderer [37]. We used AP₅₀, AP with an IoU threshold of 0.5, as the primary metric for evaluating the detection accuracy under different attacks, considering that it is a widely-used and practical metric for object detection [29, 38].

Training recipe of PBCAT. We trained the object detector with PBCAT on MS-COCO. Unless otherwise specified, the

Method	Clean (Inria)	AdvPatch	Clean (Synthetic)	AdvTexture	AdvCaT
Vanilla	96.2	37.3	86.4	0.2	0.3
LGS [35]	95.9	24.1	86.4	3.9	4.3
SAC [32]	96.2	57.1	85.4	0.3	0.6
EPGF [60]	95.1	43.2	86.7	2.9	0.4
Jedi [44]	92.3	64.4	88.1	2.3	0.7
FNC [54]	96.8	53.0	81.9	6.0	5.8
APE [21]	96.2	47.9	81.9	0.8	0.3
UDF [55]	69.1	19.3	84.9	2.2	5.8
PatchZero [53]	96.2	38.5	79.4	0.0	0.2
NAPGuard [49]	96.1	47.0	81.1	2.2	0.4
AD-YOLO [19]	90.0	39.0	90.2	0.5	0.2
l_∞ -AT (MTD) [56]	92.3	54.6	89.4	1.4	34.5
l_∞ -AT (AdvOD) [29]	95.9	56.1	92.5	30.5	39.6
PBCAT (Ours)	95.4	77.6	92.5	60.2	56.4

Table 1. The detection accuracies (AP_{50}) of models with different defense methods under adaptive attacks. Clean (Inria) and AdvPatch [45] were evaluated on the Inria dataset. Clean (Synthetic), AdvTexture [16], and AdvCaT [17] were evaluated on the synthetic dataset.




Figure 3. Visualization of examples of four attacks used in our evaluation. PGDPatch was optimized on MS-COCO. AdvPatch was optimized on the Inria dataset. AdvTexture and AdvCaT were optimized on the synthetic dataset used in [17]. PGDPatch is used to evaluate the general object detection task while other attacks were used to evaluate the downstream person detection.

patches in each image were generated with a patch perturbation step size $\alpha = 8/255$, patch perturbation intensity $\beta = 64/255$, scale factor $\lambda = \sqrt{2}/5$, and the amount of sub-patches $N = 64$. The replay parameter for FreeAT was set to be $r = 8$. Specifically, each bounding box had a 50% chance of being attached to an adversarial patch to increase the detection accuracy for objects without adversarial patches (clean objects). The global perturbations were generated with a perturbation intensity $\epsilon = 4/255$. Additional training settings basically followed the recipe proposed by Li et al. [29] (see Appendix C), which resulted in the recent SOTA robustness against the l_∞ attack. Some training examples used in PBCAT are shown in Fig. 2.

Attack evaluation setups. For the general detection task on MS-COCO, our attack evaluation used the masked PGD attack to create the adversarial square patch for each bounding box, termed PGDPatch attack. Here the iteration step was set to 200, the step size was set to 2, and the hyperparameter for the patch size was set to $\lambda = 1/5\sqrt{2}$, resulting in a patch area of 1% to 5% relative to the area of the bounding box. Note that the physical implementation of the patches created by PGDPatch was relatively dif-

ficult because the tricks for physical implementation like the TV loss [42] were not used. Instead, we cared more about the security-critical person detection task, where three actual physically realizable attacks were evaluated: AdvPatch [45], AdvTexture [16], and AdvCaT [17]. The evaluation settings for these three attacks strictly followed their original configurations in the digital world: The detector trained on MS-COCO was evaluated on the downstream person detection task directly. All of the three attacks also have their implementations and evaluations in the physical world. However, validating the effectiveness of defense methods against these physically realizable attacks in the digital world is sufficient, because the attack success rates of physically realizable attacks in the digital world is typically higher than those in the real physical world [16, 45]. In real physical world, physical implementation errors and differences in physical conditions (*e.g.*, illumination) decrease the attack success rates. Thus, if a defense method performs well in defending against these physically realizable attacks in the digital world, it can exhibit stronger robustness in the real physical world. Similar evaluation paradigms have also been adopted by many defense methods [32, 50, 60].

4.2. Robustness against Adversarial Attacks

We first compared PBCAT with the recent SOTA l_∞ -AT method for object detection, denoted as *AdvOD* [29], whose training strategy was also adopted in our work, on the general object detection task. The results shown in Appendix D.1 indicate that PBCAT achieved an AP_{50} of 37.8% averaged across all object categories and an AP_{50} of 34.5% for the person category under the PGDPatch attack, outperforming [29] by 6.1% and 4.4%, respectively. These results demonstrate the potential of PBCAT in defending against physically realizable attacks in general object detection. We then turn to the actual physically realistic attacks on the security-critical person detection task.

In the person detection task, we compared PBCAT with various defense approaches against patch-based attacks, such as input preprocessing-based methods and outlier feature filter-based methods (see Sec. 2.3). These defense methods were applied to the standardly trained detector, denoted as *Vanilla*. We also compared PBCAT with more l_∞ -AT methods, including *AdvOD* [29] and *MTD* [56]. We give the implementation details of these baseline methods in Appendix D.2. Please note that early patch-based AT works [34, 36, 50] were not evaluated as these methods were proposed originally for the classification task and it is challenging to adapt to object detection (see Sec. 3.3). All of these defense methods were evaluated in the white-box adaptive setting to show the worst-case robustness. Note that PBCAT is an AT method that has full gradient available, so we employed the direct adaptive attack rather than specific techniques like BPDA [2].

The results of different defense methods are shown in Tab. 1. These non-AT defense methods are all vulnerable to the adaptive patch attack. For the AdvPatch attack, the best non-AT defense method was SAC [32], achieving 57.1% AP_{50} . Moreover, against the stronger adaptive adversarial texture attacks, all these non-AT defense methods were broken. Interestingly, we observed that l_∞ -AT outperformed all non-AT defense methods, although the latter is explicitly designed for physically realizable attacks. Note that Li et al. [29] did not evaluate their method against physically realizable attacks in their original work, as their focus was on defending against l_∞ -bounded perturbations. Nevertheless, our results show that l_∞ -AT also enhanced robustness against physically realizable attacks. Compared with l_∞ -AT, PBCAT further improved the robustness. In particular, against AdvTexture, PBCAT improved the detection accuracy by 29.7%. Appendix D.3 visualizes some detection results of the model trained with PBCAT under physically realizable attacks.

Despite the distinct features between training examples (Fig. 2) and the evaluated attack examples (Fig. 3), PBCAT model had strong robustness against various attacks, showing its good transferability across different attacks. To

Patch	Global	Gradient	AdvPatch	AdvTexture	AdvCaT
✓			35.4	1.6	0.8
✓	✓		72.8	24.9	19.5
✓	✓	✓	77.6	63.3	56.4

Table 2. The detection accuracies (AP_{50}) of models trained with different ablation settings. AdvPatch [45] was evaluated on Inria; AdvTexture [16] and AdvCaT [17] were evaluated on the synthetic dataset used in Hu et al. [17].

Method	AdvPatch	AdvTexture	AdvCaT
Random	59.0	14.2	47.0
Gradient	77.6	63.3	56.4

Table 3. The selection strategies for sub-patches. “Random” represents selecting half of the sub-patches randomly. “Gradient” selects the sub-patches with the top 50% highest gradient values.

further validate this, we evaluated two transfer-based patch attacks [15, 18] on our model, and the results detailed in Appendix D.4 show that these transfer attacks almost lose the ability to fool the detectors trained with PBCAT.

4.3. Ablation Study

4.3.1. Effectiveness of Each Design in PBCAT

We first conducted an ablation study to show the effectiveness of each design in PBCAT: using the small area patch perturbations, using the global imperceptible noise perturbations, and using the gradient-guided patch partition and selection technique, denoted as “Patch”, “Global”, and “Gradient”, respectively. Using only “patch” indicates that its shape is a square. The results are shown in Tab. 2. We can see that all of these designs contribute to enhancing robustness against physically realizable attacks. By incorporating the global perturbations in conjunction with adversarial patches, the robust accuracy against the AdvPatch [45] attack increased to 72.8%. Introducing the gradient-guided patch partitioning and selection technique further improved robustness across all attacks.

Additionally, we examined different sub-patch selection strategies, with results shown in Tab. 3. “Random” refers to selecting half of the sub-patches randomly, and “Gradient” selects the top 50% with the highest gradient values. Notably, selecting sub-patches randomly after partitioning resulted in sub-optimal robustness, further highlighting the importance of gradient guidance in our selection process.

4.3.2. Influence of Hyper-Parameters

The number of sub-patches. Since the sampled large patch is partitioned into sub-patches, an exploration is required to identify the optimal number of partitions. We conducted three experiments by varying the number of sub-patches N : 16 (4×4), 64 (8×8), and pixel-level partitioning (each pixel as a sub-patch). Other settings followed those

Sub-patches	AdvPatch	AdvTexture	AdvCaT
16	78.3	50.8	46.2
64	77.6	60.2	56.4
Pixel-level	67.4	20.4	59.4

Table 4. AP₅₀ of models trained with different numbers of sub-patches. Pixel-level means each pixel is treated as a sub-patch.

Scale factor λ	AdvPatch	AdvTexture	AdvCaT
$2\sqrt{2}/10$	78.3	50.8	46.2
$3\sqrt{2}/10$	80.4	49.1	38.5
$4\sqrt{2}/10$	63.4	24.6	43.6

Table 5. AP₅₀ of models trained with different sizes of sampled patches. Here the number of the sub-patches was set to 16.

described in Sec. 4.1. Tab. 4 demonstrates that the number of sub-patches requires a balance.

The size of the sampled patch. λ controlled the size of the sampled patch, representing the ratio of the edge length of a square patch to the diagonal length of the bounding box. We consider three models with different λ values to examine their effect. The results shown in Tab. 5 indicate that simply enlarging the size of the sampled patch was ineffective when defending against large-area texture attacks. Moreover, too large sampled patches had negative effects even when defending against the patch-based attack.

The selection of top- k values. In Sec. 3.3 we select the top 50% of sub-patches with the highest norm values. Here we conducted experiments with varying top- k values. The results in Appendix D.5 show that selecting the top 50% strikes a favorable balance, yielding robust performance across different scenarios.

4.4. Effectiveness across Object Detectors

PBCAT requires no assumption about the structure of the detector and has shown success on the two-stage Faster R-CNN detector in the above experiments. Here we evaluated PBCAT with two additional detectors to validate its effectiveness on more models. Here we used FCOS [46], a typical single-stage object detector, DN-DETR [25], a transformer-based detector.

The detectors were trained on MS-COCO. Three methods were evaluated for each detector: standard training, AdvOD, and PBCAT. The patch perturbation step size was $\alpha = 4/255$. Patch perturbation intensity was $\beta = 32/255$ and patch size λ was 0.2. Other settings basically followed those on Faster R-CNN. The detectors with standard training was taken from the `mmdetection` [5] repository directly. The evaluation results on FCOS, DN-DETR are shown in Tab. 6, Tab. 7, respectively. FCOS trained by PBCAT achieved an AP₅₀ of 58.0% on the Inria dataset under the AdvPatch attack, surpassing l_∞ -AT by 28.1%. Its AP₅₀ on AdvTexture and AdvCaT are also much higher than that

Method	PGDPatch	AdvPatch	AdvTexture	AdvCaT
Vanilla	17.6	53.9	0.0	0.1
l_∞ -AT (AdvOD)	26.7	29.9	30.1	17.7
PBCAT (Ours)	27.8	58.0	55.1	26.0

Table 6. The detection accuracies (AP₅₀) of the FCOS models trained with different methods against attacks on different tasks.

Method	PGDPatch	AdvPatch	AdvTexture	AdvCaT
Vanilla	7.3	2.4	0.3	2.8
l_∞ -AT (AdvOD)	32.0	23.5	0.0	0.1
PBCAT (Ours)	32.7	56.3	16.8	56.8

Table 7. The detection accuracies (AP₅₀) of the DN-DETR models trained with different methods against attacks on different tasks.

in Vanilla and l_∞ -AT. Similarly, PBCAT significantly improved the robustness of DN-DETR. These results show the effectiveness of PBCAT across diverse object detectors.

4.5. Robustness in the Physical World

To further confirm the robustness of PBCAT in real-world scenarios, we evaluated the detector by using four videos provided by the authors of AdvTexture [16] and AdvCaT [17]. The videos feature two types of texture-based adversarial clothing, tested in both indoor and outdoor environments. A demonstration video in the Supplementary Materials, compares the detection performance of Faster R-CNN without defense mechanisms and with BCAT. The results indicate that PBCAT significantly improves robustness against physical-world adversarial attacks. Specifically, while the vanilla Faster R-CNN often fails to detect persons wearing adversarial clothing, the model trained with PBCAT successfully detects them in most frames.

5. Discussion and Conclusion

In this work, we introduce PBCAT, a novel adversarial training method to defend against physically realizable attacks. With extensive experiments on the security-critical person detection task, we demonstrated the effectiveness of PBCAT across various scenarios under strong adaptive attacks. Notably, PBCAT demonstrates significant improvements over previous SOTA l_∞ -AT method when defending against the AdvTexture [16] attack. We encourage future work in enhancing adversarial robustness to consider a broader range of attacks beyond l_∞ attack and patch-based attacks. Additionally, our work highlights that AT is still one of the most promising ways to achieve robustness against physically realizable attacks. The social impact of this work is discussed in Appendix E.

Limitation. Similar to most AT works [29, 58], PBCAT slightly decreased clean accuracy. It is an open question whether there is an internal trade-off between robustness against physically realizable attacks and clean accuracy.

Acknowledgement

This work was supported by the National Natural Science Foundation of China (No. U2341228).

References

- [1] Eduardo Arnold, Omar Y. Al-Jarrah, Mehrdad Dianati, Saber Fallah, David Oxtoby, and Alex Mouzakitis. A survey on 3d object detection methods for autonomous driving applications. *IEEE Trans. Intell. Transp. Syst.*, 20(10):3782–3795, 2019. 1
- [2] Anish Athalye, Nicholas Carlini, and David A. Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *Int. Conf. Mach. Learn. (ICML)*, pages 274–283, 2018. 2, 3, 7
- [3] Tom B. Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. *arXiv preprint arXiv:1712.09665*, 2017. 1, 5
- [4] Nicholas Carlini and David A. Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy*, pages 39–57, 2017. 2
- [5] Kai Chen, Jiaqi Wang, Jiangmiao Pang, et al. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 8
- [6] Pin-Chun Chen, Bo-Han Kung, and Jun-Cheng Chen. Class-aware robust adversarial training for object detection. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 10420–10429, 2021. 2
- [7] Ping-yeh Chiang, Michael J. Curry, Ahmed Abdelkader, Aounon Kumar, John Dickerson, and Tom Goldstein. Detection as regression: Certified object detection with median smoothing. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2020. 3
- [8] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *Int. Conf. Mach. Learn. (ICML)*, pages 1310–1320, 2019. 3
- [9] Francesco Croce and Matthias Hein. Minimally distorted adversarial examples with a fast adaptive boundary attack. In *Int. Conf. Mach. Learn. (ICML)*, pages 2196–2205, 2020. 2
- [10] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 886–893, 2005. 2, 5
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 248–255. IEEE Computer Society, 2009. 12
- [12] Ziyi Dong, Pengxu Wei, and Liang Lin. Adversarially-aware robust object detector. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 297–313, 2022. 2
- [13] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *Int. Conf. Learn. Represent. (ICLR)*, 2015. 2
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 770–778, 2016. 5
- [15] Yu-Chih-Tuan Hu, Jun-Cheng Chen, Bo-Han Kung, Kai-Lung Hua, and Daniel Stanley Tan. Naturalistic physical adversarial patch for object detectors. In *Int. Conf. Comput. Vis. (ICCV)*, pages 7828–7837, 2021. 7, 13
- [16] Zhanhao Hu, Siyuan Huang, Xiaopei Zhu, Fuchun Sun, Bo Zhang, and Xiaolin Hu. Adversarial texture for fooling person detectors in the physical world. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 13297–13306. IEEE, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [17] Zhanhao Hu, Wenda Chu, Xiaopei Zhu, Hui Zhang, Bo Zhang, and Xiaolin Hu. Physically realizable natural-looking clothing textures evade person detectors via 3d modeling. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 16975–16984. IEEE, 2023. 1, 2, 3, 4, 5, 6, 7, 8
- [18] Hao Huang, Ziyang Chen, Huanran Chen, Yongtao Wang, and Kevin Zhang. T-SEA: transfer-based self-ensemble attack on object detection. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 20514–20523, 2023. 7, 13
- [19] Nan Ji, Yanfei Feng, Haidong Xie, Xueshuang Xiang, and Naijin Liu. Adversarial YOLO: defense human detection patch attacks via detecting adversarial patches. *arXiv preprint arXiv:2103.08860*, 2021. 6, 13
- [20] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. Ultralytics yolov8, 2023. 5
- [21] Taeheon Kim, Youngjoon Yu, and Yong Man Ro. Defending physical adversarial attack on object detection via adversarial patch-feature energy. In *ACM Int. Conf. Multimedia*, pages 1905–1913. ACM, 2022. 1, 3, 6, 13
- [22] Chethan Kumar, R Punitha, et al. Yolov3 and yolov4: Multiple object detection for surveillance applications. In *International Conference on Smart Systems and Inventive Technology (ICSSIT)*, pages 1316–1321, 2020. 1
- [23] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *ICLR (Workshop)*. OpenReview.net, 2017. 12
- [24] Cassidy Laidlaw, Sahil Singla, and Soheil Feizi. Perceptual adversarial robustness: Defense against unseen threat models. In *Int. Conf. Learn. Represent. (ICLR)*, 2021. 3
- [25] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M. Ni, and Lei Zhang. DN-DETR: accelerate DETR training by introducing query denoising. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 13609–13617, 2022. 1, 5, 8, 14
- [26] Xiao Li, Ziqi Wang, Bo Zhang, Fuchun Sun, and Xiaolin Hu. Recognizing object by components with human prior knowledge enhances adversarial robustness of deep neural networks. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 45(7):8861–8873, 2023. 2
- [27] Xiao Li, Yining Liu, Na Dong, Sitian Qin, and Xiaolin Hu. Partimagenet++ dataset: Scaling up part-based models for robust recognition. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 396–414, 2024. 5
- [28] Xiao Li, Wei Zhang, Yining Liu, Zhanhao Hu, Bo Zhang, and Xiaolin Hu. Language-driven anchors for zero-shot adversarial robustness. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 24686–24695, 2024. 2
- [29] Xiao Li, Hang Chen, and Xiaolin Hu. On the importance of backbone to the adversarial robustness of object detectors.

- IEEE Transactions on Information Forensics and Security*, 20:2387–2398, 2025. 3, 5, 6, 7, 8, 12, 13
- [30] Xiao Li, Wenxuan Sun, Huanran Chen, Qiongxiu Li, Yining Liu, Yingzhe He, Jie Shi, and Xiaolin Hu. Adbm: Adversarial diffusion bridge model for reliable adversarial purification. In *Int. Conf. Learn. Represent. (ICLR)*, 2025. 2
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 740–755, 2014. 2, 5
- [32] Jiang Liu, Alexander Levine, Chun Pong Lau, Rama Chelappa, and Soheil Feizi. Segment and complete: Defending object detectors against adversarial patch attacks with robust patch detection. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 14953–14962. IEEE, 2022. 1, 3, 6, 7, 13
- [33] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *Int. Conf. Learn. Represent. (ICLR)*, 2018. 2
- [34] Jan Hendrik Metzen, Nicole Finnie, and Robin Huttmacher. Meta adversarial training against universal patches. *Int. Conf. Mach. Learn. (ICML)*, 2021. 1, 2, 3, 5, 7, 12
- [35] Muzammal Naseer, Salman H. Khan, and Fatih Porikli. Local gradients smoothing: Defense against localized adversarial attacks. In *WACV*, pages 1300–1307. IEEE, 2019. 3, 6, 13
- [36] Sukrut Rao, David Stutz, and Bernt Schiele. Adversarial training against location-optimized adversarial patches. In *ECCV Workshop*, pages 429–448. Springer, 2020. 1, 2, 3, 4, 5, 7, 12
- [37] Nikhila Ravi, Jeremy Reizenstein, David Novotný, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. arxiv preprint arxiv:2007.08501. *CoRR*, 2020. 5
- [38] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 5
- [39] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 39(6):1137–1149, 2017. 1, 2, 5
- [40] Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do adversarially robust imagenet models transfer better? In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2020. 3, 12
- [41] Ali Shafahi, Mahyar Najibi, Amin Ghiasi, Zheng Xu, John P. Dickerson, Christoph Studer, Larry S. Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, pages 3353–3364, 2019. 2, 5, 12
- [42] 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 *CCS*, pages 1528–1540. ACM, 2016. 6
- [43] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *Int. Conf. Learn. Represent. (ICLR)*, 2014. 1, 2
- [44] Bilel Tarchoun, Anouar Ben Khalifa, Mohamed Ali Mahjoub, Nael B. Abu-Ghazaleh, and Ihsen Alouani. Jedi: Entropy-based localization and removal of adversarial patches. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 4087–4095, 2023. 3, 6, 13
- [45] Simen Thys, Wiebe Van Ranst, and Toon Goedemé. Fooling automated surveillance cameras: Adversarial patches to attack person detection. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 49–55, 2019. 1, 3, 4, 5, 6, 7
- [46] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: fully convolutional one-stage object detection. In *Int. Conf. Comput. Vis. (ICCV)*, pages 9626–9635, 2019. 1, 2, 5, 8, 14
- [47] Florian Tramèr, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On adaptive attacks to adversarial example defenses. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2020. 2, 3
- [48] Zeyu Wang, Xianhang Li, Hongru Zhu, and Cihang Xie. Revisiting adversarial training at scale. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 24675–24685, 2024. 5
- [49] Siyang Wu, Jiakai Wang, Jiejie Zhao, Yazhe Wang, and Xianglong Liu. Nappguard: Towards detecting naturalistic adversarial patches. In *CVPR*, pages 24367–24376. IEEE, 2024. 6, 13
- [50] Tong Wu, Liang Tong, and Yevgeniy Vorobeychik. Defending against physically realizable attacks on image classification. In *Int. Conf. Learn. Represent. (ICLR)*, 2020. 1, 2, 3, 4, 5, 6, 7, 12
- [51] Zuxuan Wu, Ser-Nam Lim, Larry S Davis, and Tom Goldstein. Making an invisibility cloak: Real world adversarial attacks on object detectors. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 1–17. Springer, 2020. 1
- [52] 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 *Eur. Conf. Comput. Vis. (ECCV)*, pages 665–681, 2020. 1
- [53] Ke Xu, Yao Xiao, Zhaoheng Zheng, Kaijie Cai, and Ram Nevatia. Patchzero: Defending against adversarial patch attacks by detecting and zeroing the patch. In *WACV*, pages 4621–4630. IEEE, 2023. 6, 13
- [54] Cheng Yu, Jiansheng Chen, Youze Xue, Yuyang Liu, Weitao Wan, Jiayu Bao, and Huimin Ma. Defending against universal adversarial patches by clipping feature norms. In *Int. Conf. Comput. Vis. (ICCV)*, pages 16414–16422. IEEE, 2021. 1, 3, 6, 13
- [55] Youngjoon Yu, Hong Joo Lee, Hakmin Lee, and Yong Man Ro. Defending person detection against adversarial patch attack by using universal defensive frame. *IEEE Trans. Image Process.*, 31:6976–6990, 2022. 1, 3, 6, 13
- [56] Haichao Zhang and Jianyu Wang. Towards adversarially robust object detection. In *Int. Conf. Comput. Vis. (ICCV)*, pages 421–430, 2019. 2, 6, 7
- [57] Huan Zhang, Hongge Chen, Zhao Song, Duane S. Boning, Inderjit S. Dhillon, and Cho-Jui Hsieh. The limitations of adversarial training and the blind-spot attack. In *Int. Conf. Learn. Represent. (ICLR)*, 2019. 3
- [58] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P. Xing, Laurent El Ghaoui, and Michael I. Jordan. Theoretically

- principled trade-off between robustness and accuracy. In *Int. Conf. Mach. Learn. (ICML)*, pages 7472–7482, 2019. [2](#), [8](#)
- [59] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu. Object detection with deep learning: A review. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, 30(11):3212–3232, 2019. [1](#), [2](#)
- [60] Guangzhi Zhou, Hongchao Gao, Peng Chen, Jin Liu, Jiao Dai, Jizhong Han, and Ruixuan Li. Information distribution based defense against physical attacks on object detection. In *Int. Conf. Multimedia and Expo*, pages 1–6. IEEE, 2020. [1](#), [3](#), [6](#), [13](#)