To Make Yourself Invisible with Adversarial Semantic Contours

Yichi Zhang[a,1], Zijian Zhu[b,1], Hang Su[a], Jun Zhu[a,*], Shibao Zheng[b], Yuan He[c], Hui Xue[c]

[a]*Department of Computer Science and Technology, Institute for Artificial Intelligence, THBI Lab, Tsinghua University, Beijing 100084, China*
[b]*Institute of Image Communication and Network Engineering, Shanghai Jiao Tong University, Shanghai 200240, China*
[c]*Alibaba Group, Hangzhou 311121, China*

ABSTRACT

Modern object detectors are vulnerable to adversarial examples, which may bring risks to real-world applications. The sparse attack is an important task which, compared with the popular adversarial perturbation on the whole image, needs to select the potential pixels that is generally regularized by an ℓ_0-norm constraint, and simultaneously optimize the corresponding texture. The non-differentiability of ℓ_0 norm brings challenges and many works on attacking object detection adopted manually-designed patterns to address them, which are meaningless and independent of objects, and therefore lead to relatively poor attack performance. In this paper, we propose Adversarial Semantic Contour (ASC), an MAP estimate of a Bayesian formulation of sparse attack with a deceived prior of object contour. The object contour prior effectively reduces the search space of pixel selection and improves the attack by introducing more semantic bias. Extensive experiments demonstrate that ASC can corrupt the prediction of 9 modern detectors with different architectures (*e.g.*, one-stage, two-stage and Transformer) by modifying fewer than 5% of the pixels of the object area in COCO in white-box scenario and around 10% of those in black-box scenario. We further extend the attack to datasets for autonomous driving systems to verify the effectiveness. We conclude with cautions about contour being the common weakness of object detectors with various architecture and the care needed in applying them in safety-sensitive scenarios.

1. Introduction

Deep Neural Networks (DNNs) have obtained remarkable success on object detection (Ren et al., 2015; He et al., 2017; Liu et al., 2016; Redmon and Farhadi, 2018; Carion et al., 2020), which involves simultaneous object localization and classification. Such methods have been widely used in security-sensitive scenarios, including video surveillance, facial recognition and autonomous driving systems. However, recent research has shown the vulnerability of DNN-based models (Szegedy et al., 2013; Goodfellow et al., 2014; Madry et al., 2017; Dong et al., 2018), and modern detectors are vulnerable to carefully designed adversarial perturbations (Carlini and Wagner, 2017a; Lu et al., 2017; Xie et al., 2017; Shen et al., 2019; Wu et al., 2020), which raises safety concerns in real-

(a) Object Detection   (b) Prior Contour   (c) Invisible with ASC

Fig. 1: We conduct a sparse attack on object detectors guided by a contour prior, which can be rough and outlined manually. With pixel sampling and texture optimization, we successfully make the woman in the green bounding box invisible to DETR (Carion et al., 2020), while other objects remain detectable, by modifying only 2% of the pixels in the object area.

world applications. Therefore, it is crucial to investigate the adversarial weakness of object detection, which can then be addressed to strengthen the robustness of object detectors and further improve safety in scenarios such as auto-driving.

Methods for object detection have much more diverse

---

*Corresponding author

e-mail: zyc22@mails.tsinghua.edu.cn (Yichi Zhang), dcszj@mail.tsinghua.edu.cn (Jun Zhu)

[1]Equal Contribution

pipelines compared to those for image classification, which makes it nontrivial to conduct adversarial attacks within a generic framework. Specifically, modern detectors can be classified into one-stage detectors (*e.g.*, Yolo (Redmon et al., 2016; Redmon and Farhadi, 2016, 2018)), which use a single forward pass to make predictions, and two-stage detectors (*e.g.*, R-CNN (Girshick et al., 2014), Faster R-CNN(Ren et al., 2015)), which generate region proposals by Region Proposal Network (RPN) and conduct classification on them. RPN brings with redundant region proposals and smaller feature maps of Region of Interest (RoI), which may neutralize adversarial perturbations and lead to more challenges for adversarial attack on two-stage detectors. The rising Transformer architecture (Parmar et al., 2018; Guo et al., 2021; Liu et al., 2021) is being applied to object detectors, either as backbone (Liu et al., 2021) or as detection head (*e.g.*, DETR and its incremental versions (Carion et al., 2020; Zhu et al., 2020b; Liu et al., 2022a)). Some studies reveal that it's more difficult to conduct adversarial attacks on Vision Transformers (ViT), because they learn more generalized contextual information (Shao et al., 2021), but the adversarial robustness of Transformer-based detectors still remain rarely studied.

There has been extensive research on generating adversarial examples for object detection. Early methods like DAG (Xie et al., 2017) and UEA (Wei et al., 2018) perturb the whole digital image with imperceptible noise bounded by ℓ_∞ or ℓ_2 norm. Recently, sparse attack has drawn an increasing amount of attention (Liu et al., 2018; Wu et al., 2020; Thys et al., 2019; Brown et al., 2017; Huang et al., 2020). It explores the robustness of neural network from a different setting, where a constraint on ℓ_0 norm of the perturbation, *i.e.*, the number of pixels to be modified, is imposed, instead of limiting the perturbation budget. It's equally worthy as global attack to be studied in both algorithms and applications, because it's a typical NP-hard problem and its setting is closer to that in real-world attack, where only limited area of the environment can be disturbed.

However, this is a problem of constrained optimization and also technically more challenging to solve as it requires to optimize on the texture and select attacked pixels simultaneously, and the pixel selection part is generally a combinatorial problem that is sparse and non-differentiable. Therefore, we have to resort to find approximate solutions. Traditional ℓ_0 attack methods in image classification (Carlini and Wagner, 2017b; Croce and Hein, 2019) usually disturb scattered pixels over the whole image, which are selected with heuristic or other strategies. Apart from those, many recent studies on sparse attack on object detectors focus more on perturbing areas which meet certain connected patterns that are commonly manually predefined and independent of the object properties (Huang et al., 2020; Saha et al., 2019; Thys et al., 2019; Liu et al., 2018; Wu et al., 2020; Bao, 2020). This is also known as patch-based attack. From another perspective, these predesigned patterns serve as prior knowledge for pixel selection. However, the meaningless shape and location of the prior patterns are not able to fully capture the object characteristics and can sacrifice the effectiveness in terms of the attack success rate or the ℓ_0 norm and the efficiency of texture optimization.

To address the discussed problem, we present Adversarial Semantic Contour (ASC) to improve the sparse attack performance on object detectors. We define ASC as an Maximum A Posteriori (MAP) estimate under a Bayesian formulation of ℓ_0-constrained sparse attack and design a prior over the object contours for pixel selection. ASC can effectively attack various object detectors in an efficient manner. In particular, our prior knowledge of object contour comes from the insights of previous works that reveal the importance of contour and shape for object detection. In (Alexe et al., 2010), the authors thoroughly analyze the question "What is an object?", and point out that an object should have closed boundaries and a different appearance against the background. This implies that the object contours carry significant information about the object. Additionally, it has shown that object detectors may benefit from shape cues (Geirhos et al., 2018). Based on the informative prior contour, we can optimize the texture with a gradient method similar to PGD (Madry et al., 2017) technically. To further uplift the attack performance, we model a prior distribution defined by the contour and adopt MAP estimation to select the pixels to be perturbed via sampling. The introduction of contour fits the characteristics of the task of object detection, and significantly enhance the attack.

We conduct comprehensive experiments to verify the attack effectiveness and efficiency of our method. We first apply ASC in the popular dataset COCO (Lin et al., 2014) and in auto-driving with two different datasets of CityScapes (Cordts et al., 2016) and BDD100K (Yu et al., 2020) under the task of object vanishing. Experimental results show that, with the same ℓ_0 norm constraints (*e.g.*, no more than 5% of the object area in the COCO case), we can achieve better attack both than the popular predesigned patterns and than the traditional ℓ_0 attack methods with further optimization in pixel selection. The convergence curves and transfer-based attacks indicate that ASC achieves a successful attack with fewer iterations of texture optimization and better transferability in black-box adversarial attacks thanks to the appropriate contour priors. Additional studies also prove the power of our methods in different tasks and on defense models. Our results suggest that semantic contour is a common weakness to various detectors by testing on 9 different modern detectors.

In summary, our contributions are:

- We introduce a simple yet effective prior of object contours by considering the characteristics of an object detector, which promotes attack effectiveness and efficiency along with transferability as a consequence.

- We propose a novel Bayesian framework for adversarial attacks on object detection, which provides a generic solution for critical sparse attacks on object detection with ℓ_0 constraints.

- Comprehensive experiments are conducted on nine modern detectors (including one-stage, two-stage detectors and Detection Transformers) within COCO and two auto-driving datasets, and the results verify the effectiveness, efficiency and transferability of our method, which reveal the common weakness of object detectors and indicate the

potential safety hazards of object detection and its application, *e.g.*, auto-driving systems.

2. Related Work

In this section, we briefly review previous methods for object detection and adversarial attacks on that.

2.1. Object Detection

In "What is an object?" (Alexe et al., 2010), three distinctive characteristics of objects are investigated, including "a closed boundary in space", "a different appearance" and "salience against the background". Additionally, four cues for objectness measurement are applied. Multi-Scale Saliency and Color Contrast are closely related to foreground color, while Edge Density and Straddleness correlate more with object contour. An object can be detected by optimizing the parameters of the measurement in a Bayesian framework. Taking the parameter ϕ of Edge Detection for instance, the optimal ϕ^* is maximized via the posterior probability of windows covering the ground-truth labels as positives as

$$\phi^* = \arg\max_{\phi} \prod_{w \in \mathcal{W}^{\text{obj}}} p_{\phi}(\text{obj}|\text{ED}(w, \phi)) \tag{1}$$
$$= \arg\max_{\phi} \prod_{w \in \mathcal{W}^{\text{obj}}} \frac{p_{\phi}(\text{ED}(w, \phi)|\text{obj}) \cdot p(\text{obj})}{\sum_{c \in \{\text{obj,bg}\}} p_{\phi}(\text{ED}(w, \phi)|c) \cdot p(c)},$$

where $\text{ED}(w, \phi)$ is the edge density of window w with parameter ϕ, $p(\text{obj})$ is the probability of the window to be predicted as an object and $\mathcal{W}^{\text{obj}}$ is the set of windows bounding a ground-truth object.

Modern DNN-based object detectors are mainly categorized into one-stage detectors, such as Yolo (Redmon et al., 2016; Redmon and Farhadi, 2016, 2018) and SSD (Liu et al., 2016), which make predictions with one forward pass, and two-stage detectors, such as Faster R-CNN (Ren et al., 2015) and Mask R-CNN (He et al., 2017), which are mostly based on R-CNN (Girshick et al., 2014). In addition to ResNet (He et al., 2016) and VGGNet (Simonyan and Zisserman, 2014), recent Vision Transformers (ViT) like Swin-T (Liu et al., 2021) can also be used as backbone network to extract features. Besides, Detection Transformer (DETR) (Carion et al., 2020) applies Transformer (Vaswani et al., 2017) architecture and removes postprocessing operations such as NMS. The adversarial robustness of ViTs compared to CNNs is under heated discussion (Mahmood et al., 2021; Bai et al., 2021; Shao et al., 2021) and that of Transformer-based detectors still remains further study.

2.2. Adversarial Attack on Object Detection

Adversarial examples for DNNs were first put forward in (Szegedy et al., 2013). DNN-based object detectors have been proven to be vulnerable to adversarial examples (Carlini and Wagner, 2017a; Lu et al., 2017; Xie et al., 2017; Wei et al., 2018). The generation of adversarial perturbations is often a constrained optimization problem and existing methods can be categorized according to their ℓ_p constraints. Some methods apply ℓ_∞ as the bound for perturbations to create noise imperceptible to humans, such as DAG (Xie et al., 2017) and UEA (Wei

et al., 2018). Nevertheless, these methods add perturbations over the whole image, which can be impractical in a real-world attack. Compared to the ℓ_∞ attack, the ℓ_0 attack, also known as sparse attack, which aims to modify a limited number of pixels, is also important to be studied.

Sparse attack needs to select the attacked pixels and optimize the texture simultaneously, while the pixel selection is non-differentiable. In the work of C&W (Carlini and Wagner, 2017b), an ℓ_0 attack pipeline using an ℓ_2 adversary is proposed, whose pixel selection tightly relies on target models and takes many rounds of attacks to get the solution. PGD_0 (Croce and Hein, 2019) selects a given number of pixels to modify in every iteration of texture optimization. However, these methods are proposed in the scenario of classification and always modified scattered pixels. (Brown et al., 2017) first introduces the adversarial patch, which has become popular for object detection (Liu et al., 2018; Thys et al., 2019; Saha et al., 2019; Wu et al., 2020). Manually designed patterns, such as square patches (Liu et al., 2018; Thys et al., 2019; Saha et al., 2019) and grids (Wu et al., 2020), are adopted to avoid time-costly search and need only iterative optimization for their textures. However, the meaningless patterns, which are pre-designed with certain shapes irrelevant to the object characteristics, limit the attack effectiveness in terms of the attack success rate or the ℓ_0 cost and the convergence efficiency for texture optimization.

3. Methodology

In this section, we present our formulation of sparse attack on an object detector constrained by ℓ_0 norm within a Bayesian framework and introduce a prior of object contour for the pixel selection.

In general, to attack with ℓ_0 constraints, we need to decide which pixels are to be disturbed and their resulting adversarial textures. We consider a formulation of finding the optimal adversarial noise for each image in a general form as

$$\delta^* = \arg\max_{\delta} J(g_\omega(x + \delta), y), \text{ s.t. } \ell_0(\delta) \leq N_0, \tag{2}$$

where x is the original image, y is its corresponding ground-truth label, g_ω denotes the detector with parameter ω representing the structure and weights of the model, $J(g_\omega(\cdot), y)$ denotes the objective function that describes the similarity between the detector output and the ground-truth label y, and δ is the adversarial noise. We maximize the objective function under the constraint that the ℓ_0 norm of δ is smaller than a constant N_0. Based on Eq. (1), the detector prediction is a likelihood given the parameters and input. In this perspective, we can rewrite our formulation of an ℓ_0 attack as an Maximum Likelihood estimation (MLE) with an ℓ_0 constraint:

$$\delta^* = \arg\max_{\delta} \log p_\omega(\bar{y}|x + \delta), \text{ s.t. } \ell_0(\delta) \leq N_0, \tag{3}$$

where we substitute maximizing the objective function J with maximizing the log-likelihood of the wrong prediction of $\bar{y}$.

One challenge of problem (3) is that δ involves pixel selection and texture optimization simultaneously, which are coupled with each other. We propose to divide it into two parts as

$\delta = t \odot M$ and alternately optimize them. Here, $t \in \mathbb{R}^{h \times w \times 3}$ is the texture, $M \in \{0, 1\}^{h \times w}$ is the binary mask for pixel selection and $\odot$ is element-wise multiplication. Then, the optimization of δ can be decoupled as

$$\max_{M,t} \log p_\omega(\bar{y}|x + t \odot M)$$
$$= \max_{M,t} \log p_\omega(\bar{y}|M, t), \ \text{s.t.} \ \ell_0(M) \leq N_0, \tag{4}$$

where the ℓ_0 constraint only concerns M and we ignore the fixed input x in the condition of likelihood for notational clarity.

The challenge of this problem is that the optimization of M with an ℓ_0 constraint is non-differentiable and NP-hard in general. We may get a good approximation of the ℓ_0 norm using ℓ_1 norm, but the non-differentiable modules in diverse object detectors (Girshick et al., 2014) still make the problem intractable only with gradient methods. We here adopt appropriate prior knowledge to better approximate the optimization, motivated by the claim in (Juan et al., 2020) that introducing non-uniform randomization priors can cope effectively with non-smooth optimization. We introduce a prior distribution for $\delta = (M, t)$ which can guide the optimization in a Bayesian framework with MAP estimation constrained by ℓ_0 norm as

$$M^*, t^* = \arg\max_{M,t} \log \left(p_\omega(\bar{y}|M, t) p(M, t) \right)$$
$$= \arg\max_{M,t} \log \left(p_\omega(\bar{y}|M, t) p(t|M) p(M) \right), \tag{5}$$
$$\text{s.t.} \ \ell_0(M) \leq N_0,$$

where $p(M, t)$ is the joint prior distribution for $\delta = (M, t)$ based on prior knowledge, and it can be factorized as $p(t|M)p(M)$.

The above formulation subsumes some existing methods that solve Eq. (3). For instance, the methods using heuristic strategies (such as C&W-ℓ_0 (Carlini and Wagner, 2017b) and PGD$_0$(Croce and Hein, 2019)) have $p(M, t)$ as a uniform distribution. As for sparse patch-based methods, such as fixed pre-designed patterns like squares and grids (Liu et al., 2018; Thys et al., 2019; Wu et al., 2020), their prior distribution for $p(M)$ can be seen as an impulse function and they set $p(t|M)$ uniformly with no special assumptions because texture t can be optimized based on gradient methods given a certain binary mask M. However, these patterns lack object characteristics and information, and cannot achieve successful attacks effectively on various detectors. Similarly, we also focus on $p(M)$ in our work but with a more appropriate prior, which can facilitate the optimization, and the problem then becomes

$$M^*, t^* = \arg\max_{M,t} \log \left(p_\omega(\bar{y}|M, t) p(M) \right),$$
$$\text{s.t.} \ \ell_0(M) \leq N_0. \tag{6}$$

Below, we introduce an object contour as a more reasonable prior for M, which carries object semantics and improves the attack effectiveness, and explain the implementation details of contour acquisition, texture optimization and the MAP estimation based on the contour to further improves the attack performance.
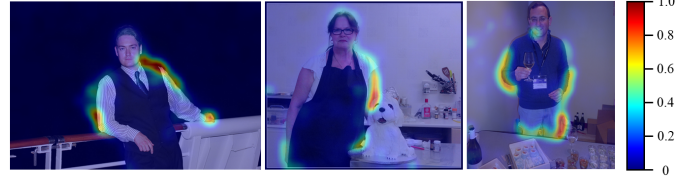

Fig. 2: We generate heat-maps of three images on DETR by measuring the normalized adversarial contribution by replacing different areas with adversarial noise, which indicates that the object contour is more important in adversarial attack, compared with the inside/outside regions.

3.1. Using Object Contour?

In this section, we conduct a thorough analysis and provide evidence that the object contour is not only a semantic but an informative area of the object itself, which can facilitate adversarial attacks on various object detectors.

In (Alexe et al., 2010), Alexe *et al*. found distinctive attributes of an object and proposed different metrics closely related to color and boundary. In the analysis of shape bias and texture bias (Geirhos et al., 2018), the authors found that training a Shape-ResNet with Stylized ImageNet (SIN), which maintains shape bias while the texture is corrupted, can improve the performance in object detection. This indicates that object detectors can benefit more from shape bias when localizing the objects. Besides these studies of training better models from shape cues, we further explore the significance of contour for object detection under adversarial perturbations. We define a metric of normalized Adversarial Contribution (nAC), which measures the relative contribution of an area during adversarial attack compared with other areas in an image:

$$t_a = \arg\max_t \log p_\omega(\bar{y}|x + t \odot M_a), \tag{7}$$

$$\text{AC}(a) = \log p_\omega(\bar{y}|x + t_a \odot M_a) - \log p_\omega(\bar{y}|x), \tag{8}$$

$$\text{nAC}(a) = \frac{\text{AC}(a) - \min_a \text{AC}}{\max_a \text{AC} - \min_a \text{AC}}, \tag{9}$$

where a represents an area on the image, and M_a denotes the binary mask for area a. Then, we add adversarial perturbations t_a to the area and evaluates $\text{AC}(a)$ by getting the difference of the model predictions caused by the added adversarial noise in a. And $\text{nAC}(a)$ indicates the normalized contribution among the areas across the whole image. We evaluate different contribution of an area a at different area of "Inside", "Contour" or "Outside" corresponding to an object.

In Fig. 2, we show the representative heat-maps describing the normalized adversarial contribution on DETR. This demonstrates that perturbations on color and texture inside the object may not be sufficient to pose any threat to object detectors with the ℓ_0 constraint. Moreover, after examining 500 images on 9 detectors, we compare the normalized adversarial contribution of three classes of area quantitatively in Tab. 1. The results show that areas around the boundary indeed contribute more during adversarial attack, which implies that, by disturbing around the object boundary with same number of pixels, the characteristics of the object may be corrupted with higher probabilities compared with disturbing inside or outside the object.

Table 1: Normalized Adversarial Contribution (↑, %) of area examined.

	Outside	Inside	Contour
SSD	7.64	23.58	**25.75**
YOLOv3	8.67	30.72	**31.05**
YOLOX	9.16	18.48	**21.72**
Faster R-CNN	11.49	23.46	**26.77**
Mask R-CNN	11.61	22.76	**27.00**
Mask R-CNN Swin-T	10.06	17.72	**22.68**
DETR	9.38	11.95	**14.85**
DAB-DETR	11.65	15.98	**19.03**
Deformable DETR	9.33	19.64	**23.95**

3.2. Contour Acquisition

Here, we introduce some implementation details of our method on how to acquire the appropriate prior contour. The object semantic contour is a intuitive concept, which describes the outline of an object, which is composed of the boundary between the object and its background, and to be more precise, those between different parts of the object (*e.g.*, those between legs, arms, head and body for human targets). Therefore, the contour acquisition can be done by intuitive approaches, like human annotations, when we are attacking a few images. As for batches of images, we can obtain the object contour from the result of instance segmentation (He et al., 2017), or part segmentation (Lin et al., 2020) if we need the contour to be more detailed. Practically, by eroding the segmentation mask, we can easily acquire sparse but highly-connected contour.

3.3. Problem Formulation

In this section, we present our problem formulation of texture optimization based on the contour and the MAP estimation to further improve the attack performance.

Texture Optimization. Based on the acquired prior contour, represented as M, we keep it fixed and adopt a method similar to PGD (Madry et al., 2017) to optimize the texture, which projects the updated texture t into the acceptable range for pixel values as

$$t \leftarrow \text{Clip}_{[0,1]}(t + \alpha \cdot \nabla_t \log p_\omega(\bar{y}|M, t)), \quad (10)$$

where we only optimize t with gradients of objective function for a given M and α is a hyper-parameter as step size. For an image with pixel values ranging from 0 to 1, we clip the modified pixel values within the acceptable range after every updating step.

Further Optimization for Pixel Selection. The fixed mask M is likely to limit the attack performance and the further optimization can be achieved by adjusting the contour with better pixel selection. With an appropriate prior for M, using MAP estimation is a natural deduction from the Bayesian framework of ℓ_0-constrained sparse attack. In particular, we may parameterize the connected contour with θ, which corresponds to its binary mask representation $M(\theta)$, and obeys a prior normal distribution $\pi(\theta)$. The problem in Eq. (6) can be further reformulated as a problem parameterized by θ. Then the optimal solution to the adversarial noise of (θ^*, t^*) can be solved via MAP
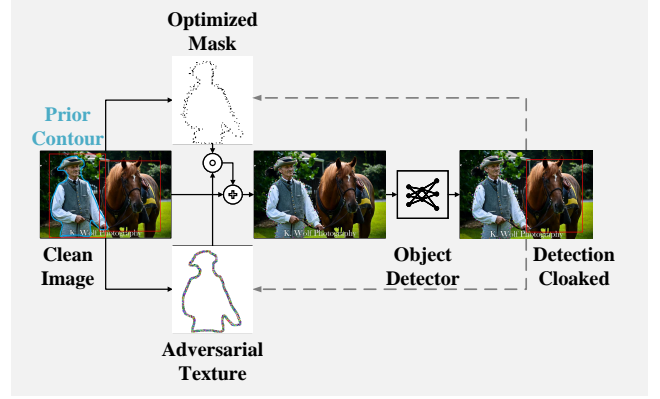

Fig. 3: We separate the perturbation δ into texture t and binary mask M. By introducing a prior distribution of M around the contour, we can easily achieve a successful attack on object detectors with a MAP estimation under the Bayesian framework.

Algorithm 1 Adversarial Semantic Contour

Input: inference model with weights of ω; raw image x; label y; hyper-parameters α, β.
Output: adversarial example x'
1: Get the prior contour C_p parameterized by θ_0
2: Initial $\theta = \theta_0$ and t
3: **while** attack fails **do**
4:     Update θ as
        $\theta \leftarrow (1 - \beta)\theta + \beta(\nabla_\theta \log p_\omega(\bar{y}|\theta, t) + \theta_0)$
5:     Optimize t for θ iteratively as
        $t \leftarrow \text{Clip}_{[0,1]}(t + \alpha \cdot \nabla_t \log p_\omega(\bar{y}|\theta, t))$
6: **end while**
7: Project the optimized θ to binary mask M
8: **return** $x' \leftarrow x + t \odot M$

---

estimation as

$$\theta^*, t^* = \underset{\theta, t}{\arg\max} \log p_\omega(\bar{y}|\theta, t) + \log \pi(\theta),$$
$$\text{s. t. } \ell_0(M(\theta)) \leq N_0, \quad (11)$$

where $\theta \sim \mathcal{N}(\theta_0, \Sigma)$ with θ_0 and Σ respectively being the prior contour and an identity matrix. Based on the prior contour with parameter of θ_0, we update θ iteratively using the gradient w.r.t. θ of the expression in Eq. (11) with θ_0 as an initialization as

$$\theta \leftarrow (1 - \beta)\theta + \beta(\nabla_\theta \log p_\omega(\bar{y}|\theta, t) + \theta_0), \quad (12)$$

where β is a hyper-parameter that can be controlled to avoid $l_0(M(\theta))$ from exceeding the ℓ_0 constraint and t is the optimized texture given θ.

However, the optimization is still intractable because the process involves the projection from continuous θ to sparse M. In practice, to achieve better attack successful rate, we take Monte Carlo sampling around the prior contour, which results in scattered pixels generated around the prior contour. We summarize our framework in Algorithm 1 and in Fig. 3 graphically.
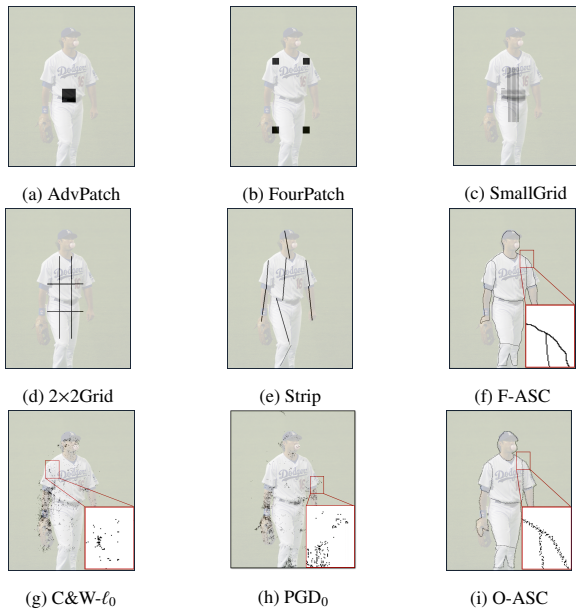
Fig. 4: Examples of mask patterns for different methods.

(a) AdvPatch (b) FourPatch (c) SmallGrid

(d) 2×2Grid (e) Strip (f) F-ASC

(g) C&W-ℓ_0 (h) PGD$_0$ (i) O-ASC

(a) AdvPatch 32.67% (b) FourPatch 25.11% (c) SmallGrid 6.18%

(d) 2×2Grid 4.11% (e) Strip 9.88% (f) F-ASC 2.94%

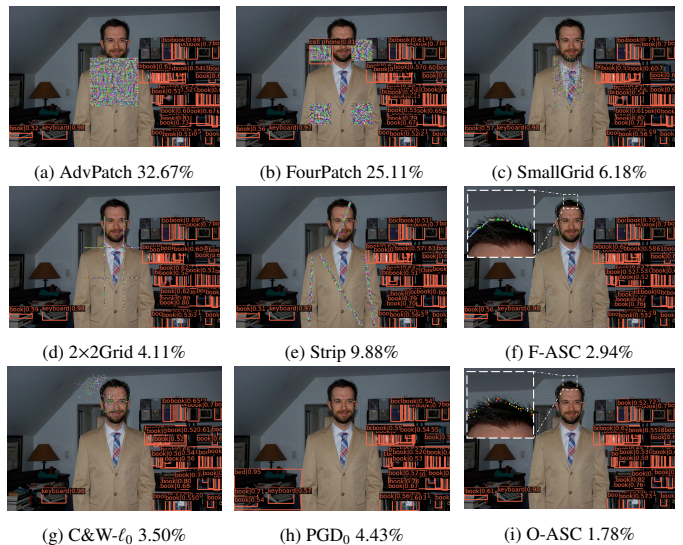(g) C&W-ℓ_0 3.50% (h) PGD$_0$ 4.43% (i) O-ASC 1.78%

Fig. 5: Examples of object vanishing on COCO. Our method manipulate the fewest pixels to make a person invisible to DETR. The percentage indicates the ratio of modified pixel number to pixels of the object area.

4. Experiments

4.1. Experiment Settings

Datasets. We use three datasets, including COCO (Lin et al., 2014) as a universal dataset for object detection and Cityscapes (Cordts et al., 2016) along with BDD100K (Yu et al., 2020) as two typical datasets for autonomous driving systems.

Models. We mainly attack 9 models in total: 3 one-stage detectors including SSD512 (Liu et al., 2016), Yolov3 (Redmon and Farhadi, 2018) and YoloX (Ge et al., 2021), 3 two-stage detectors including Faster R-CNN (FRCN) (Ren et al., 2015), Mask R-CNN (MRCN) (He et al., 2017), both with backbone of ResNet50 (He et al., 2016) and Mask R-CNN with Swin Transformer (Swin-T) (Liu et al., 2021) as backbone (MRCN-Swin), and 3 Detection Transformers including DETR (Carion et al., 2020), Deformable DETR (Def-DETR) (Zhu et al., 2020b) and DAB-DETR (Liu et al., 2022a). These models are all trained on COCO by mmdetection (Chen et al., 2019), except for DAB-DETR[2], whose official code is newly released.

Methods. Our method (ASC) is based on the semantic boundaries of objects. "F-ASC" stands for the fixed case that directly uses prior contours, which are acquired from part-segmentation (Lin et al., 2020) for "Person" in COCO and ground-truth segmentation for other objects. Because we stress the ℓ_0 constraints, we have the area cost of our contour pattern bounded by 5% of the object area in COCO and 10% in the other two datasets. We show the example patterns in Fig. 4 (f) and (i). To verify the performance of our contour-based attack, we adopt five prior patterns as baseline methods to compare in our experiments. The five prior patterns are mainly in three classes, which are square patches ("AdvPatch" (Thys

et al., 2019) and "FourPatch" in Fig. 4 (a) and (b)), grids ("2×2Grid" (Wu et al., 2020) and "SmallGrid" in Fig. 4 (c), (d)) and "Strip" in Fig. 4 (e) generated based on part segmentation. We restrict the area cost of these methods to be the same as our contour pattern. Note that all methods including our proposed ASC optimize the texture via gradient-based method, which involves iterative forward inference and backward propagation, the contour acquisition for ASC with segmentation model only needs one forward inference and brings negligible extra computation under our unified attack framework.

4.2. Attack Effectiveness

In this section, we focus on the task of object vanishing, *i.e.*, to make the target object undetected, and verify the effectiveness of our method in terms of Successful Detection Rate (SDR), *i.e.*, if a prediction with an IoU above 0.5 with the target object and a confidence score above a threshold (0.3 for Deformable DETR and DAB-DETR, 0.5 for the rest), then we take it as a success. We make the objects invisible to the detectors by maximizing the possibility of the object being the background, which is the most risky for the application of object detection. To boost the performance of our method, we take "O-ASC" for the optimized case that consists of the further pixel sampling for examples against which the attacks fail, as mentioned in Sec.3.2. Since it results in "scattered pixels" and the baseline patterns are all connected, to have the comparison fairer, we take two more baselines, C&W (Carlini and Wagner, 2017b) and PGD$_0$ (Croce and Hein, 2019), which are traditional ℓ_0 attack methods proposed for image classification and generates perturbations over the whole image. The proposed ℓ_0 attack in C&W, whose binary mask is shown in Fig. 4 (g), follows the pipeline that one pixel is removed in every iteration and we increase the step size for pixels for efficiency. As for PGD$_0$, which is in Fig. 4 (h), we take the settings in its original paper and its application in object detection (Li and Zhao, 2020).

---

[2]https://github.com/IDEA-opensource/DAB-DETR

Table 2: Successful Detection Rate (%, ↓) in object vanishing for 'Person' in COCO.

	SSD512	Yolov3	YoloX	FRCN	MRCN	MRCN-Swin	DETR	DAB-DETR	Def-DETR
Clean	95.6	92.1	99.4	98.3	98.9	99.6	99.6	98.5	97.3
AdvPatch	6.2	11.5	25.2	85.0	88.7	81.7	94.5	74.7	67.1
FourPatch	22.8	17.0	43.3	76.4	77.7	79.5	90.4	67.1	46.2
SmallGrid	2.0	0.4	5.6	46.9	53.3	51.5	85.7	28.1	20.0
2×2Grid	0.9	**0.2**	5.3	28.3	32.9	41.0	75.3	18.5	7.8
Strip	1.1	0.8	3.6	16.8	21.4	45.3	47.9	9.9	4.4
F-ASC (ours)	**0.3**	0.7	**1.7**	**9.7**	**10.8**	**24.4**	**25.3**	**4.0**	**2.6**
C&W-ℓ_0	0.2	0.6	2.3	13.1	15.1	29.4	46.6	5.4	1.8
PGD$_0$	0.4	0.3	8.4	8.8	9.3	22.4	21.8	6.6	3.1
O-ASC (ours)	**0.1**	**0.1**	**0.2**	**2.0**	**2.7**	**10.7**	**5.2**	**1.0**	**0.3**

Table 3: Successful Detection Rate (%, ↓) in object vanishing for other categories in COCO.

	SSD512	Yolov3	YoloX	FRCN	MRCN	MRCN-Swin	DETR.	DAB-DETR	Def-DETR
Clean	90.1	89.4	96.2	90.7	92.0	97.8	97.9	93.3	93.2
AdvPatch	5.7	13.2	26.1	57.3	62.6	68.8	71.8.	45.6	47.4
FourPatch	12.2	12.4	30.5	45.3	48.8	69.5	63.3	38.0	26.6
SmallGrid	3.0	1.9	6.6	20.8	22.5	31.2	53.1.	14.5	11.8
2×2Grid	**2.8**	**1.0**	4.2	12.1	13.0	28.9	45.0	9.7	6.9
Strip	6.5	4.0	13.5	25.7	26.7	44.8	47.5	16.2	9.9
F-ASC (ours)	3.5	2.9	**4.0**	**9.4**	**9.2**	**25.2**	**18.8**	**5.3**	**2.0**
C&W-ℓ_0	**1.0**	**0.4**	1.6	8.7	8.0	22.1	29.8	3.0	1.2
PGD$_0$	2.7	1.0	5.0	5.1	7.0	19.0	18.0	4.1	3.4
O-ASC (ours)	1.8	1.0	**0.7**	**3.8**	**3.5**	**13.5**	**5.6**	**0.9**	**0.7**

Experiments on COCO. COCO (Lin et al., 2014) is a popular dataset with 80 categories for object detection and other visual tasks. We randomly select 1000 images with labels of "Person" and 1000 of other categories as the attack targets.

To better visualize the margins of attack performance within the methods as in Fig. 5, we examine the minimum ℓ_0 cost to achieve successful attacks against DETR (Carion et al., 2020) for each method, by increasing l_0 constraint at certain step size until the person is cloaked. This has a different setting from the following quantitative experiments with same l_0 budget, but their conclusions of effectiveness are consistent. Our methods, both "F-ASC" with fixed contour and "O-ASC" with optimized selection, require far fewer pixels to make the person invisible, indicating the attack effectiveness of our method. It is noted that compared to "AdvPatch," which is the most popular pattern, our methods reduce the number of modified pixels from around 30% to less than 3%.

Quantitative experimental results are shown in Tab. 2 and Tab. 3, which are for "Person" and other categories in COCO respectively. We list the results of these two types of objects separately because the prior contour for the "Person" is acquired from part segmentation, which is more precise and detailed, while that for the other is only from ground-truth segmentation label. Both tables indicate the similar trends. We see that for one-stage detectors, their adversarial robustness are the weakest and most of the methods can lead to quite low SDRs. Our methods achieve lower or similar SDRs on these detectors, both for "F-ASC" compared to the other five patterns and for "O-ASC"

compared to the two traditional ℓ_0 attack methods. As for the other two types of detectors, two-stage detectors and Detection Transformers, our methods can lead to the lowest SDRs with quite significant gaps. To be specific, the gaps of SDR between "F-ASC" and other prior patterns on DETR are more than 22.6% for "Person" and 26.2% for other categories, and those between "O-ASC" and two ℓ_0 attack methods on it are more than 16.6% for "Person" and 12.4% for other categories. The reason is that the object contour serves as a common feature of the object among the detectors and carries more information of the object existence, which can mislead various detectors. Also, with further sampling, "O-ASC" is much more effective than "F-ASC". For instance, on DETR, the declines in SDRs are 10.1% and 13.2% respectively. This indicates the update of pixel selection based on the prior contour can be promising.

Two noticeable phenomena appear in our results. The first one is that, with Swin-T as the backbone network, the adversarial robustness of Mask R-CNN is improved compared to that with ResNet50. The resulting SDRs of almost all methods become higher, indicating that the Swin-T backbone help Mask R-CNN resist the adversarial attack. It has been studied that ViT models have global receptive fields, which allows to model global context and therefore lead to robustness of learned features (Naseer et al., 2021). Swin Transformer, though it computes self-attention within local windows to improve efficiency, still have larger receptive fields (window size of at least 7×7) than ResNet50 (kernel size of 3×3) (Liu et al., 2022b) and results in better robustness to our attack. The second one is that,

Table 4: Successful Detection Rate (%, ↓) in Cityscapes.

	SSD512	Yolov3	YoloX	FRCN	MRCN	MRCN-Swin	DETR	DAB-DETR	Def-DETR
Clean	99.3	93.6	99.9	100.0	100.0	100.0	99.9	99.9	99.9
AdvPatch	15.8	31.1	41.1	80.1	82.6	88.3	90.4	80.7	64.5
FourPatch	13.8	24.1	43.1	57.0	61.6	85.9	79.0	47.6	30.3
SmallGrid	5.2	8.5	4.1	39.6	40.8	60.7	78.1	20.8	20.6
2×2Grid	3.1	6.4	4.8	22.9	31.6	63.8	71.8	5.8	6.8
Strip	3.1	11.8	14.6	30.4	37.9	72.4	63.5	13.7	6.9
F-ASC(ours)	**2.3**	**3.0**	**3.5**	**21.4**	**30.5**	**55.6**	**22.4**	**1.6**	**0.2**
C&W-ℓ_0	**0.1**	0.6	2.1	4.7	7.6	38.6	36.7	1.4	1.3
PGD$_0$	5.3	1.1	0.4	11.5	13.1	37.4	31.4	2.6	1.8
O-ASC(ours)	**0.1**	**0.2**	**0.2**	**2.9**	**6.1**	**26.3**	**2.1**	**0.2**	**0.0**

Table 5: Successful Detection Rate (%, ↓) in BDD100K.

	SSD512	Yolov3	YoloX	FRCN	MRCN	MRCN-Swin	DETR	DAB-DETR	Def-DETR
Clean	93.4	83.1	98.8	98.1	98.5	99.2	98.3	97.0	97.5
AdvPatch	23.4	21.7	49.8	79.6	81.1	81.9	91.1	77.2	66.3
FourPatch	24.8	8.3	41.4	61.2	63.9	78.9	77.8	51.7	34.7
SmallGrid	4.0	2.2	9.0	39.6	42.5	52.8	77.5	19.6	13.4
2×2Grid	2.8	1.0	9.6	17.5	20.1	51.4	64.6	5.9	3.9
Strip	6.8	1.9	17.5	24.1	26.4	58.4	59.8	12.2	5.4
F-ASC(ours)	**2.7**	**0.8**	**6.5**	**14.8**	**18.9**	**51.1**	**29.6**	**2.8**	**1.0**
C&W-ℓ_0	**0.0**	**0.0**	1.3	7.5	10.7	33.5	38.6	1.7	1.4
PGD$_0$	4.6	0.6	2.0	8.7	10.4	37.7	43.7	4.3	4.3
O-ASC(ours)	0.7	**0.0**	**1.0**	**6.5**	**8.7**	**31.8**	**8.0**	**1.1**	**0.4**

among the three Detection Transformers, the two incremental versions of DETR appear to be much easier to be attacked, which suggests their improvement measures to uplift the prediction accuracy in terms of mean Average Precision (mAP) bring with worse adversarial robustness. Concretely, they are DETR, DAB-DETR and Deformable DETR, arranged from the most robust to the weakest. One typical difference between DETR and its incremental versions is that DETR uses Cross Entropy Loss while the other two models use Focal Loss (Lin et al., 2017). This lead to better predicting performance but sacrifice the robustness, maybe because Focal Loss is proposed based on Binary Cross Entropy Loss to address the imbalance of positive and negative samples, and to put more weights on objects that are harder to distinguish. Different sizes of receptive fields may also contribute to this phenomenon. DETR (Carion et al., 2020) adopts cross-attention mechanism whose receptive field covers the whole feature map, while Deformable DETR (Zhu et al., 2020b) employs deformable attention with sparse sampling points which significantly decrease the size of receptive field. As for DAB-DETR (Liu et al., 2022a), it uses the scale information of dynamic anchor boxes to modulate the attention, similar to soft RoI pooling, which leads to its receptive field smaller than DETR but larger than Deformable DETR. The rank of the equivalent receptive fields corresponds to that of their robustness to our attack.

Moreover, the separate results of "Person" and other categories in the two tables again confirm the correctness of our claim about object semantics. For meaningless pre-designed patterns, C&W-ℓ_0 and PGD$_0$, most of their resulting SDRs for

(a) Successful Detection   (b) Invisible

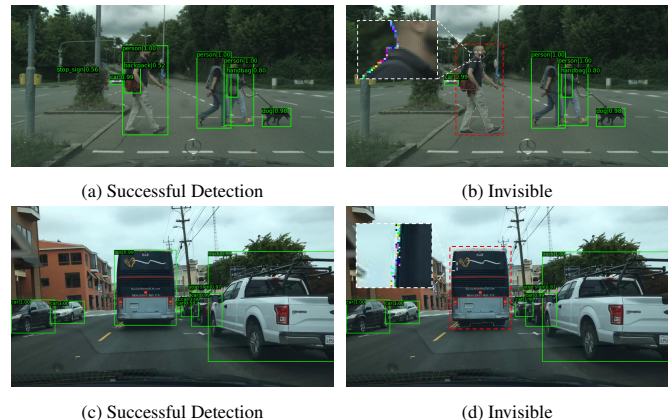(c) Successful Detection   (d) Invisible

Fig. 6: Examples of our attack on Cityscapes and BDD100K. (a) and (b) are images from Cityscapes, (c) and (d) are images from BDD100K. The red boxes with dash line highlight that the objects are invisible to the detectors.

"Person" are higher than those for other categories. This indicates that detectors tend to be more accurate and robust when detecting humans. However, this kind of gap is much smaller for "Strip" and ASC. Since "Strip" and object contour are acquired from part segmentation for "Person", which is more precise and thereby carries more accurate semantic information, they can still achieve similar or even better attack performance than that of other categories.

Experiments on Auto-Driving Datasets. Modern detectors have shown their power in scene analysis and are applied in many scenarios, such as auto-driving. We test the at-
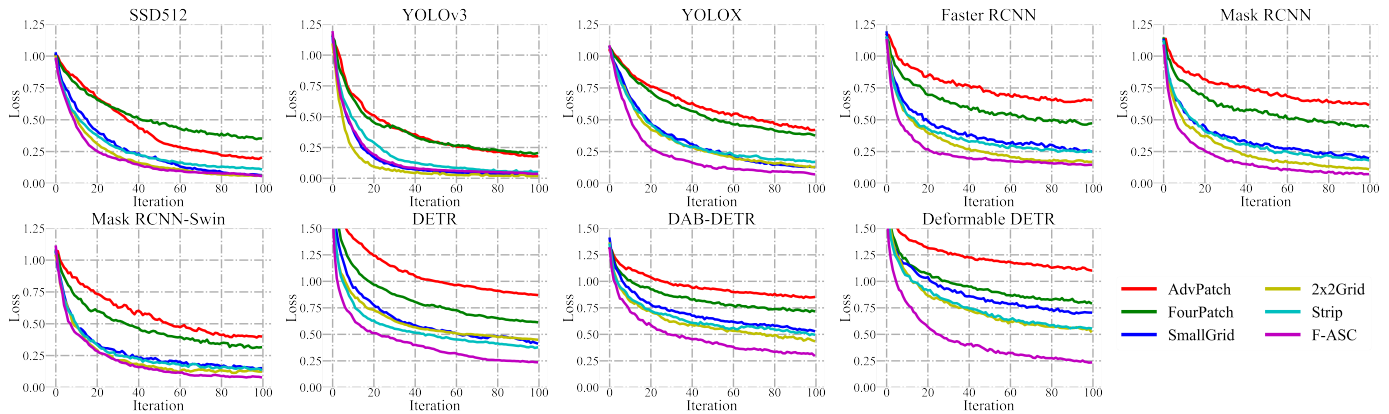
Fig. 7: Convergence Curve on different types of models in the task of object detection in COCO.

Table 6: Successful Detection Rate (%,↓) for transfer-based black-box adversarial attacks across various detectors.

Ensemble models		MRCN	DAB-DETR	Retina Net	Corner Net	Cascade R-CNN	SCNet	Query Inst	TOOD	PAA	SABL RetinaNet	DetectoRS	Dynamic R-CNN	Auto Assign	Point Rend	YOLACT
	Clean	95.7	96.5	90.1	96.5	95.3	96.6	87	90.9	98.5	92.3	98.6	94.2	97.4	97.5	90.2
FRCN +YOLO +SSD	AdvPatch	90.1	94.4	83.2	94.1	91	93.5	82.1	82.5	95.9	89.1	97.6	88.4	95.9	96.1	86.1
	FourPatch	87.5	93.5	80.5	89.8	86.8	90.3	74.4	81.7	95.3	87.6	97.9	83.9	94	94.6	85
	SmallGrid	80.3	91.7	68.2	88.8	82.4	89.5	71.7	68.8	91.4	79.8	95.8	77.5	91.9	92.4	81
	2×2Grid	75.4	88.4	66.7	81.8	77.3	86.3	66.6	63.6	88.4	77.8	96	73.5	89	91.7	75.5
	Strip	77.4	88.2	69	77.2	74.7	85.9	62.4	71.4	90.5	79.5	95.9	69.6	89.9	90.1	76.1
	F-ASC	**63.7**	**79.4**	**52.8**	80.7	**65.9**	**79.4**	**53.5**	**59**	**82.4**	**70.9**	94.4	**55.9**	**84.7**	**87.9**	72.7
Swin-MRCN +YOLOX +DETR	AdvPatch	95	94.8	88	94.4	94.5	95.3	81.7	86.1	97.1	90.1	98	93.1	96.1	96.6	88.2
	FourPatch	94.6	94.7	87.3	91.6	93.5	94.9	78.1	86.1	97.4	89	98.2	91.1	95	95.7	87.6
	SmallGrid	92.7	92	81.7	90.5	91.9	93.4	72.1	77.2	95.6	81.4	96.4	91.6	92.8	94.7	84.6
	2×2Grid	89.9	88	80	84.4	88.1	91.3	68.1	74.5	94	81.5	96.6	86.3	91.4	93.2	79
	Strip	91.3	90	83	83.7	89.3	91.4	71.8	79.7	95.3	83.3	97	88.2	93.1	93.9	82.6
	F-ASC	**85.7**	**82.3**	**71.5**	**83.6**	**83.7**	**87.5**	**62.8**	**70.9**	**90.4**	**76.4**	**95.5**	**81.5**	**89**	**92**	**78.3**

tack effectiveness of our method using auto-driving datasets of Cityscapes (Cordts et al., 2016) and BDD100K (Yu et al., 2020) toward object detection, specifically with object vanishing, to demonstrate that the fragility of DNN-based models should draw more caution to their application in safety-sensitive fields. We select 1000 objects of persons or vehicles, which are the most common objects in auto-driving, as the attack targets from each dataset respectively. As shown in Fig. 6, the pedestrian and bus in the front are cloaked by ASC, which uncovers the maximum vulnerability of modern detectors in the auto-driving scenarios.

Tab. 4 and Tab. 5 shows the results of attacks in Cityscapes and BDD100K. "F-ASC" outperforms other prior patterns on all detectors and "O-ASC" outperforms the scattered ℓ_0 attack methods on most of the detectors excluding SSD512. The trends of results are similar with those in COCO (Lin et al., 2014) and suggests that the phenomena are statistically consistent. The results show that our method not only works effectively in COCO, but also demonstrates huge safety-risks to the auto-driving scenarios.

4.3. Study on Efficiency

In this section, we analyze the efficiency for different priors. We compare the convergence speed during the texture optimization of different prior patterns. We randomly select 100 images that can be successfully attacked by each fixed prior and draw the convergence curves of their average loss in the task of object

vanishing with COCO in Fig. 7. We see the convergence speed of "F-ASC" is the fastest for 8 out of 9 detectors, except for Yolov3. The relative orders of convergence are similar to the attack effectiveness. The result indicates that the contour prior has stronger information and better efficiency when conducting attacks compared to other manually designed and meaningless priors. The faster convergence also implies the fact that modern detectors of different architecture are all sensitive to the area of object contour.

4.4. Transfer-based Attack in Black-box Scenarios

In order to demonstrate the benefits of a proper prior for the transfer-based black-box adversarial attack among different models, we use 1000 images from COCO and generate the adversarial examples for object vanishing on two sets of ensemble models respectively. One ensemble is Faster RCNN, SSD512 and YOLOv3, which are classic DNN-based detectors, while the other ensemble is Mask RCNN-Swin, YOLOX and DETR, which are with the latest and distinctive architectures. The transfer adversarial examples modify around 15% pixels of the object area.

The results in Tab. 6 show that "F-ASC" with object characteristics, has achieve the lowest SDRs on 15 black models, compared to other fixed prior patterns. Again, this proves that object contour prior shares the common characteristics of detectors with different architectures (Mask R-CNN (MRCN) (He et al., 2017), DAB-DETR (Liu et al., 2022a), RetinaNet (Lin

Table 7: Successful Detection Rate (%, ↓) in object mislabeling.

	SSD512	Yolov3	YoloX	FRCN	MRCN	MRCN-Swin	DETR	DAB-DETR	Def-DETR
Clean	94.6	91.1	99.4	97.7	98.3	99.5	99.4	98.3	96.9
AdvPatch	4.4	9.8	31.7	82.0	85.3	73.0	91.8	73.6	66.1
FourPatch	19.7	11.5	43.2	71.6	73.6	66.1	88.7	66.7	45.0
SmallGrid	1.7	0.8	6.3	42.4	48.9	38.2	81.0	28.6	16.5
2×2Grid	0.5	**0.2**	4.6	24.7	27.5	31.7	70.8	16.2	8.0
Strip	0.9	0.3	3.2	14.5	18.9	31.2	41.8	9.2	4.7
0F-ASC	**0.2**	0.6	**1.9**	**7.4**	**8.9**	**16.1**	**19.8**	**5.0**	**2.5**
C&W-ℓ_0	**0.0**	0.2	1.1	9.3	10.8	17.6	42.7	3.9	1.7
PGD$_0$	0.5	0.5	11.2	7.1	9.8	16.6	14.2	7.0	3.1
O-ASC	**0.0**	**0.0**	**0.0**	**1.2**	**2.2**	**3.9**	**4.9**	**0.9**	**0.4**

Table 8: CIoU Distance (↓) (Zheng et al., 2020) in bounding box shift.

	SSD512	Yolov3	YoloX	FRCN	MRCN	MRCN-Swin	DETR	DAB-DETR	Def-DETR
Clean	0.8416	0.8239	0.9358	0.8712	0.8882	0.9125	0.9358	0.9231	0.9103
AdvPatch	0.0660	0.1595	0.3942	0.7045	0.7372	0.6929	0.8851	0.7317	0.6877
FourPatch	0.2050	0.1539	0.4318	0.6336	0.6644	0.6822	0.8482	0.6544	0.4757
SmallGrid	0.0299	0.0210	0.1123	0.4055	0.4576	0.4625	0.8193	0.3443	0.2280
2×2Grid	**0.0123**	**0.0100**	0.0862	0.2705	0.2907	0.3947	0.7441	0.2159	0.1073
Strip	0.0234	0.0249	0.0730	0.1895	0.2108	0.4138	0.5044	0.1079	0.0633
F-ASC	0.0170	0.0201	**0.0343**	**0.0965**	**0.1088**	**0.2337**	**0.2964**	**0.0368**	**0.0262**
C&W-ℓ_0	0.0134	0.0175	0.0277	0.1527	0.1633	0.2857	0.4700	0.0480	0.0385
PGD$_0$	0.0089	0.0067	0.1793	0.0814	0.0922	0.1423	0.1783	0.0612	0.0232
O-ASC	**0.0057**	**0.0078**	**0.0109**	**0.0536**	**0.0601**	**0.1349**	**0.1092**	**0.0122**	**0.0093**

et al., 2017), CornerNet (Law and Deng, 2018), Cascade R-CNN (Cai and Vasconcelos, 2018), SCNet (Liu et al., 2020), QueryInst (Fang et al., 2021), TOOD (Feng et al., 2021), PAA (Kim and Lee, 2020), SABL RetinaNet (Wang et al., 2020), DetectoRS (Qiao et al., 2020), Dynamic R-CNN (Zhang et al., 2020), AutoAssign (Zhu et al., 2020a), PointRend (Kirillov et al., 2019), YOLACT (Bolya et al., 2019)), which can therefore degenerate their performance more effectively. The average drop of SDR among detectors with "F-ASC" compared to clean detection are 22.26% and 12.41%. The results show that the adversarial noises around object contour are more transferable, meaning that most modern detectors are adversarially sensitive and vulnerable to the adversarial noises around object contour. This further indicates a serious potential danger for real-world applications, *e.g.*, autonomous driving systems.

4.5. Additional Study

In this section, we will deliver more additional experiments to further study the performance our proposed methods. Object detection involves object localization and classification simultaneously, which makes it a multi-task problem. Therefore, the attack objectives can also be multiple. Besides object vanishing which can be taken as a targeted attack, we can also generate adversarial examples to make the detector mislabel the objects or fail in localizing the objects. By carrying out these two tasks, we may further confirm the effectiveness of our method. Additionally, as we know that adversarial attack algorithms can have different behaviors on defense models, we also test our method

on several defensive detectors. The experimental settings are consistent with those of the experiments on 1000 "Person"s in COCO.

Object Mislabeling. In this task, we carry out an untargeted attack on the detectors by minimizing the classification confidence. We also use SDR as the metric with the additional condition that the classification of the object is correct. In general, Tab. 7 shows similar trends to object vanishing task. "F-ASC" leads other prior patterns almost all detectors, especially with large gaps on two-stage detectors and Detection Transformers. "O-ASC" further reduces the SDRs with sampling and outperforms all other traditional ℓ_0 attack methods on each detector. The results demonstrate that ASC not only perturbs the prediction of object existence, but also can corrupt the classification.

Bounding Box Shift. In this task, we aim to minimize the Intersection over Union (IoU) as far as possible, which can be categorized as an attack on regression. We select Complete IoU (CIoU) (Zheng et al., 2020) as the metric. We only evaluate the positive CIoU and treat the cases with negative CIoU or no prediction as 0, since some examples may be entirely corrupted with no prediction given. Tab. 8 shows the results of our experiments for bounding box shift with different methods. "F-ASC" can outperform all other prior patterns, except on SSD512 and Yolov3, where 2×2Grid leads by a neck. After we have optimized the examples with high CIoUs, "O-ASC" has superior performance compared with other attackers on all the nine modern detectors. Though with looser constraints on noise distribution, C&W-ℓ_0 and PGD$_0$ still falls behind "O-ASC". This indi-

Table 9: Successful Detection Rate (%, ↓) with defenses on FRCN.

	Clean	AP	FP	SG	2×2G	Strip	F-ASC	C&W	PGD_0	O-ASC
Vanilla	98.3	85.0	76.4	46.9	28.3	16.8	**9.7**	13.1	8.8	**2.0**
Gaussian	98.3	85.6	74.7	46.3	27.6	17.7	**10.1**	13.7	17.4	**4.3**
Bilateral	97.4	82.8	70.5	38.6	23.6	12.5	**10.2**	16.9	22.3	**5.6**
JPEG	97.9	88.0	79.1	57.0	42.1	27.4	**20.1**	32.5	26.0	**7.9**
PCLoss	97.6	84.6	81.0	58.8	45.1	42.8	**28.5**	22.4	18.0	**15.7**

cates that our method moves the bounding boxes farther away from the ground-truth label, suggesting that the contour pattern can offer more influence on the detector not only about the object existence, but also the object location.

Defensive Detectors. There are a variety of ways to have a defense model, including input transformation (Guo et al., 2017), adversarial training (Pang et al., 2020), and improved loss functions (Li et al., 2021). Notice the fact that we want to compare different attack methods on a defense model, using adversarial training is not reasonable enough. Therefore, we mainly consider the other two different types of defense and carry out the task of object vanishing on Faster RCNN with the experimental results shown in Tab. 9. In consideration of the convenience for reading, we replicate the results of Faster R-CNN in Tab. 2 as "Vanilla" here to make comparisons. It's straightforward to consider whether frequently-used smoothing and filtering techniques can influence the sparse noises generated by methods with better performance. We first adopt Gaussian Smoothing ($\sigma = 4$) and Bilateral Filtering ($\sigma_x = \sigma_y = \sigma_r = 1.5$) for denoising as simple defense strategies and use BPDA (Athalye et al., 2018) to generate adversarial examples. From the results, we see that both techniques present some defense effectiveness for methods with more disperse noises. Still, the trends that F-ASC as fixed prior pattern outperforms other pre-defined patterns and O-ASC leads the performance over all baselines are consistent with previous results. Then, we take two methods that are more often used in adversarial defense. We trained two robust detectors using Faster RCNN including the input manipulation of JPEG (Guo et al., 2017) and an improved loss function of Probabilistically Compact Loss (PC Loss) (Li et al., 2021). The input compression by JPEG is to neutralize the influence of adversarial noises and the usage of PC Loss instead of Cross Entropy Loss is mainly to enlarge the gaps of classification probabilities and therefore strengthen the robustness. Though the defense sacrifices the performance on clean detection a little, the improvements in adversarial robustness are significant. For JPEG Compression, we adopt BPDA (Athalye et al., 2018) with all methods by using differentiable computation to approximate the compression, and as a result, "F-ASC" achieves the lowest SDR compared to the other five patterns while "O-ASC" leads to further improvement and outperforms C&W-ℓ_0 and PGD_0. As for PC Loss, we take the same attack method, and see that our methods still have better attack effectiveness. This result suggests that the proposed contour-based sparse attack works well and can be extended to the defense models.

5. Conclusion

In this paper, we propose a novel prior-guided adversarial attack on object detection, which provides a unifying view of most current sparse adversarial attacks under the Bayesian framework. To address the limitations of existing methods, especially those manually designed, we introduce a more informative prior of object contours, which are the common characteristics for various object detectors. With the prior contour appropriately harnessed, we optimize the textures, leading to a more effective and efficient adversarial attacker, and sample the pixels around the contour to further improve the attack performance and outperform the traditional ℓ_0 methods with scattered perturbations. Extensive experiments demonstrate the superior performance compared to other methods on both COCO and the auto-driving datasets. Our work concludes with the phenomenon that object contour serves as a common weakness for various modern detectors and they are more sensitive to adversarial noises in that area. This further raises concerns over the applications of these DNN-based object detectors in safety-critical systems, *e.g.*, autonomous driving systems.

References

Alexe, B., Deselaers, T., Ferrari, V., 2010. What is an object?, in: 2010 IEEE computer society conference on computer vision and pattern recognition, IEEE. pp. 73–80.

Athalye, A., Carlini, N., Wagner, D., 2018. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples, in: Proceedings of the 35th International Conference on Machine Learning, ICML 2018. URL: https://arxiv.org/abs/1802.00420.

Bai, Y., Mei, J., Yuille, A.L., Xie, C., 2021. Are transformers more robust than cnns? Advances in Neural Information Processing Systems 34.

Bao, J., 2020. Sparse adversarial attack to object detection. arXiv preprint arXiv:2012.13692 .

Bolya, D., Zhou, C., Xiao, F., Lee, Y.J., 2019. Yolact: Real-time instance segmentation, in: ICCV.

Brown, T.B., Mané, D., Roy, A., Abadi, M., Gilmer, J., 2017. Adversarial patch. arXiv preprint arXiv:1712.09665 .

Cai, Z., Vasconcelos, N., 2018. Cascade r-cnn: Delving into high quality object detection, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 6154–6162.

Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S., 2020. End-to-end object detection with transformers, in: European Conference on Computer Vision, Springer. pp. 213–229.

Carlini, N., Wagner, D., 2017a. Adversarial examples are not easily detected: Bypassing ten detection methods, in: Proceedings of the 10th ACM workshop on artificial intelligence and security, pp. 3–14.

Carlini, N., Wagner, D., 2017b. Towards evaluating the robustness of neural networks, in: 2017 ieee symposium on security and privacy (sp), IEEE. pp. 39–57.

Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., Zhang, Z., Cheng, D., Zhu, C., Cheng, T., Zhao, Q., Li, B., Lu, X., Zhu, R., Wu, Y., Dai, J., Wang, J., Shi, J., Ouyang, W., Loy, C.C., Lin, D., 2019. MMDetection: Open mmlab detection toolbox and benchmark. arXiv preprint arXiv:1906.07155 .

Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B., 2016. The cityscapes dataset for semantic urban scene understanding, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

Croce, F., Hein, M., 2019. Sparse and imperceivable adversarial attacks, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 4724–4732.

Dong, Y., Liao, F., Pang, T., Su, H., Zhu, J., Hu, X., Li, J., 2018. Boosting adversarial attacks with momentum, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 9185–9193.

Fang, Y., Yang, S., Wang, X., Li, Y., Fang, C., Shan, Y., Feng, B., Liu, W., 2021. Instances as queries, in: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pp. 6910–6919.

Feng, C., Zhong, Y., Gao, Y., Scott, M.R., Huang, W., 2021. Tood: Task-aligned one-stage object detection, in: 2021 IEEE/CVF International Conference on Computer Vision (ICCV), IEEE Computer Society. pp. 3490–3499.

Ge, Z., Liu, S., Wang, F., Li, Z., Sun, J., 2021. YOLOX: Exceeding yolo series in 2021. arXiv preprint arXiv:2107.08430 .

Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W., 2018. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. arXiv preprint arXiv:1811.12231 .

Girshick, R., Donahue, J., Darrell, T., Malik, J., 2014. Rich feature hierarchies for accurate object detection and semantic segmentation, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 580–587.

Goodfellow, I.J., Shlens, J., Szegedy, C., 2014. Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572 .

Guo, C., Rana, M., Cisse, M., Van Der Maaten, L., 2017. Countering adversarial images using input transformations. arXiv preprint arXiv:1711.00117 .

Guo, M.H., Cai, J.X., Liu, Z.N., Mu, T.J., Martin, R.R., Hu, S.M., 2021. Pct: Point cloud transformer. Computational Visual Media 7, 187–199.

He, K., Gkioxari, G., Dollár, P., Girshick, R., 2017. Mask r-cnn, in: Proceedings of the IEEE international conference on computer vision, pp. 2961–2969.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778.

Huang, L., Gao, C., Zhou, Y., Xie, C., Yuille, A.L., Zou, C., Liu, N., 2020. Universal physical camouflage attacks on object detectors, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 720–729.

Juan, A.A., Corlu, C.G., Tordecilla, R.D., de la Torre, R., Ferrer, A., 2020. On the use of biased-randomized algorithms for solving non-smooth optimization problems. Algorithms 13. URL: https://www.mdpi.com/1999-4893/13/1/8, doi:10.3390/a13010008.

Kim, K., Lee, H.S., 2020. Probabilistic anchor assignment with iou prediction for object detection, in: ECCV.

Kirillov, A., Wu, Y., He, K., Girshick, R., 2019. PointRend: Image segmentation as rendering.

Law, H., Deng, J., 2018. Cornernet: Detecting objects as paired keypoints, in: 15th European Conference on Computer Vision, ECCV 2018, Springer Verlag. pp. 765–781.

Li, H., Zhao, Y., 2020. Fool object detectors with l0-norm patch attack .

Li, X., Li, X., Pan, D., Zhu, D., 2021. Improving adversarial robustness via probabilistically compact loss with logit constraints, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 8482–8490.

Lin, K., Wang, L., Luo, K., Chen, Y., Liu, Z., Sun, M.T., 2020. Cross-domain complementary learning using pose for multi-person part segmentation. IEEE Transactions on Circuits and Systems for Video Technology .

Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P., 2017. Focal loss for dense object detection, in: Proceedings of the IEEE international conference on computer vision, pp. 2980–2988.

Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L., 2014. Microsoft coco: Common objects in context, in:

European conference on computer vision, Springer. pp. 740–755.

Liu, J.J., Hou, Q., Cheng, M.M., Wang, C., Feng, J., 2020. Improving convolutional networks with self-calibrated convolutions, in: IEEE CVPR.

Liu, S., Li, F., Zhang, H., Yang, X., Qi, X., Su, H., Zhu, J., Zhang, L., 2022a. DAB-DETR: Dynamic anchor boxes are better queries for DETR, in: International Conference on Learning Representations. URL: https://openreview.net/forum?id=oMI9PjOb9Jl.

Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C., 2016. Ssd: Single shot multibox detector. Lecture Notes in Computer Science , 21–37URL: http://dx.doi.org/10.1007/978-3-319-46448-0_2, doi:10.1007/978-3-319-46448-0_2.

Liu, X., Yang, H., Liu, Z., Song, L., Li, H., Chen, Y., 2018. Dpatch: An adversarial patch attack on object detectors. arXiv preprint arXiv:1806.02299 .

Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B., 2021. Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030 .

Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S., 2022b. A convnet for the 2020s. arXiv preprint arXiv:2201.03545 .

Lu, J., Sibai, H., Fabry, E., 2017. Adversarial examples that fool detectors. arXiv preprint arXiv:1712.02494 .

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A., 2017. Towards deep learning models resistant to adversarial attacks. arXiv preprint arXiv:1706.06083 .

Mahmood, K., Mahmood, R., Van Dijk, M., 2021. On the robustness of vision transformers to adversarial examples, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 7838–7847.

Naseer, M.M., Ranasinghe, K., Khan, S.H., Hayat, M., Shahbaz Khan, F., Yang, M.H., 2021. Intriguing properties of vision transformers. Advances in Neural Information Processing Systems 34.

Pang, T., Yang, X., Dong, Y., Su, H., Zhu, J., 2020. Bag of tricks for adversarial training. arXiv preprint arXiv:2010.00467 .

Parmar, N., Vaswani, A., Uszkoreit, J., Kaiser, L., Shazeer, N., Ku, A., Tran, D., 2018. Image transformer, in: International Conference on Machine Learning, PMLR. pp. 4055–4064.

Qiao, S., Chen, L.C., Yuille, A., 2020. Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution. arXiv preprint arXiv:2006.02334 .

Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You only look once: Unified, real-time object detection, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 779–788.

Redmon, J., Farhadi, A., 2016. Yolo9000: Better, faster, stronger. arXiv preprint arXiv:1612.08242 .

Redmon, J., Farhadi, A., 2018. Yolov3: An incremental improvement. arXiv .

Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. Advances in neural information processing systems 28.

Saha, A., Subramanya, A., Patil, K., Pirsiavash, H., 2019. Adversarial patches exploiting contextual reasoning in object detection. arXiv preprint arXiv:1910.00068 .

Shao, R., Shi, Z., Yi, J., Chen, P.Y., Hsieh, C.J., 2021. On the adversarial robustness of visual transformers. arXiv preprint arXiv:2103.15670 .

Shen, G., Mao, C., Yang, J., Ray, B., 2019. Advspade: Realistic unrestricted attacks for semantic segmentation. arXiv preprint arXiv:1910.02354 .

Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 .

Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R., 2013. Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199 .

Thys, S., Van Ranst, W., Goedemé, T., 2019. Fooling automated surveillance cameras: adversarial patches to attack person detection, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 0–0.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need, in: Advances in neural information processing systems, pp. 5998–6008.

Wang, J., Zhang, W., Cao, Y., Chen, K., Pang, J., Gong, T., Shi, J., Loy, C.C., Lin, D., 2020. Side-aware boundary localization for more precise object detection, in: ECCV.

Wei, X., Liang, S., Chen, N., Cao, X., 2018. Transferable adversarial attacks for image and video object detection. arXiv preprint arXiv:1811.12641 .

Wu, S., Dai, T., Xia, S.T., 2020. Dpattack: Diffused patch attacks against

universal object detection. arXiv preprint arXiv:2010.11679 .

Xie, C., Wang, J., Zhang, Z., Zhou, Y., Xie, L., Yuille, A., 2017. Adversarial examples for semantic segmentation and object detection, in: Proceedings of the IEEE international conference on computer vision, pp. 1369–1378.

Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T., 2020. Bdd100k: A diverse driving dataset for heterogeneous multitask learning, in: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).

Zhang, H., Chang, H., Ma, B., Wang, N., Chen, X., 2020. Dynamic R-CNN: Towards high quality object detection via dynamic training. arXiv preprint arXiv:2004.06002 .

Zheng, Z., Wang, P., Liu, W., Li, J., Ye, R., Ren, D., 2020. Distance-iou loss: Faster and better learning for bounding box regression, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 12993–13000.

Zhu, B., Wang, J., Jiang, Z., Zong, F., Liu, S., Li, Z., Sun, J., 2020a. Autoassign: Differentiable label assignment for dense object detection. arXiv preprint arXiv:2007.03496 .

Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J., 2020b. Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 .