

Class-Aware Robust Adversarial Training for Object Detection

Pin-Chun Chen^{1,2*} Bo-Han Kung² Jun-Cheng Chen²

¹Columbia University

²Research Center for Information Technology Innovation, Academia Sinica

pc2939@columbia.edu; {bhkung, pullpull}@citi.sinica.edu.tw

Abstract

Object detection is an important computer vision task with plenty of real-world applications; therefore, how to enhance its robustness against adversarial attacks has emerged as a crucial issue. However, most of the previous defense methods focused on the classification task and had few analysis in the context of the object detection task. In this work, to address the issue, we present a novel class-aware robust adversarial training paradigm for the object detection task. For a given image, the proposed approach generates an universal adversarial perturbation to simultaneously attack all the occurred objects in the image through jointly maximizing the respective loss for each object. Meanwhile, instead of normalizing the total loss with the number of objects, the proposed approach decomposes the total loss into class-wise losses and normalizes each class loss using the number of objects for the class. The adversarial training based on the class weighted loss can not only balances the influence of each class but also effectively and evenly improves the adversarial robustness of trained models for all the object classes as compared with the previous defense methods. Furthermore, with the recent development of fast adversarial training, we provide a fast version of the proposed algorithm which can be trained faster than the traditional adversarial training while keeping comparable performance. With extensive experiments on the challenging PASCAL-VOC and MS-COCO datasets, the evaluation results demonstrate that the proposed defense methods can effectively enhance the robustness of the object detection models.

1. Introduction

Due to the recent breakthroughs of deep learning, deep learning-based approaches have achieved promising performance for many computer vision tasks, such as object recognition [8], [23], [7] and object detection [6]. However, researchers found there exists potential security issues

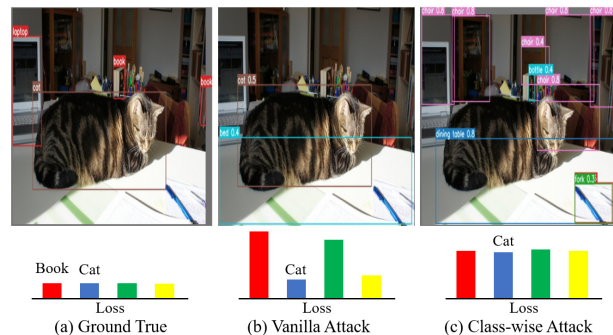


Figure 1. Detection results after attacked by different adversarial examples to the vanilla SSD model. (a) ground true, (b) we craft the adversarial example through the 10-step PGD optimization with the budget $\epsilon = 8$ on the multi-task loss as described in equation (1), (c) the detection result of the proposed class-wise attack. These detection examples show the adversarial examples generated by the proposed method can more evenly attack all the objects occurred in the image than (b).

for deep learning-based approaches. Szegedy et al. [25] first crafted adversarial examples by adding imperceptible perturbations to the input images, which can easily fool the deep learning-based classification models to generate unexpected outputs. From then on, many new attack methods, including Fast Signed Gradient Method (FGSM) [5], Deep-Fool [16], Projected Gradient Descent (PGD) [14], Carlini and Wagner Attack (C&W) [1], have been proposed to produce various adversarial examples that further expose the vulnerability of the deep learning classification models. On the other hand, object detection is one of the most important and active research fields for computer vision with plenty of real-world applications. Unfortunately, as the classification problem, it also suffers from the threat of these adversarial attacks, such as the physical adversarial patch attack to affect the steering behavior of self-driving cars [24] or the detection results of a face detector [21]. However, as compared with the development of attack methods, the defense algorithms to improve the robustness of object detection models are relatively few.

In order to defend against these attacks, various meth-

*Work done during a research assistantship at Academia Sinica.

ods have been proposed to enhance the robustness of the deep learning models, and one of the most effective defense approaches is adversarial training [26]. In addition, for the object detection task, the approaches can be roughly categorized into two types: one-stage detector [13], [18] and two-stage detector [4], [3], [19], and we focus on the one-stage detector (i.e., single-shot object detector (SSD) [13]) due to its faster detection speed and more complex nature than the two-stage detector where the nature of the two-stage detector is more similar with that of image classification task (i.e., it also performs the classification and regression tasks on the object proposals generated by the region proposal network.). Although there exists algorithms [30] to enhance the robustness of the one-stage detector, there are still some unsolved problems: vanilla adversarial training using the overall loss of one-stage object detector does not properly take all the objects occurred in an image into consideration. As shown in equation (1), the object detection loss of a specific object consists of a classification loss to identify the object class and a regression loss for bounding box regression of the object. The total loss for all the occurred objects in a given image can be written as follows:

$$\mathcal{L} = \frac{1}{N_o} \left(\sum_{i=1}^{N_o} l_{cls}(O_i, \{y_i\}, \theta) + l_{reg}(O_i, \{b_i\}, \theta) \right) \quad (1)$$

where O_i presents i -th matched default box in the image, N_o is the number of matched default boxes, l_{cls} and l_{reg} are the losses of the classification branch and regression branch respectively.

As shown in Figure 1, not all of the detected objects in an image by an object detector can be attacked successfully if we generate the adversarial examples directly using the total loss described in the equation (1) since the sub-loss for a specific object (i.e., the loss of a specific object might go to infinity.) and a specific object class (i.e., in a given image, there are more objects of a specific class than other classes.) might dominate the overall loss value during the generation process of adversarial examples. To address these issues, we present a novel class-aware robust adversarial training for the object detection task. For a given image, the proposed approach generates an universal adversarial perturbation to simultaneously attack all the occurred objects in the image through jointly maximizing the respective loss for each object. For the classification and regression losses of each object, we clip each of them respectively to avoid the situation that the specific object loss dominates the overall loss. Meanwhile, instead of normalizing the total loss with the number of objects, the proposed approach decomposes the total loss into class-wise losses and normalizes each class loss using the number of objects for the corresponding class to mitigate the situation that the loss of a specific class dominates others. The adversarial training based on the proposed class weighted loss can not only balances the influence of each class but also effectively and evenly improves adversarial robustness of trained mod-

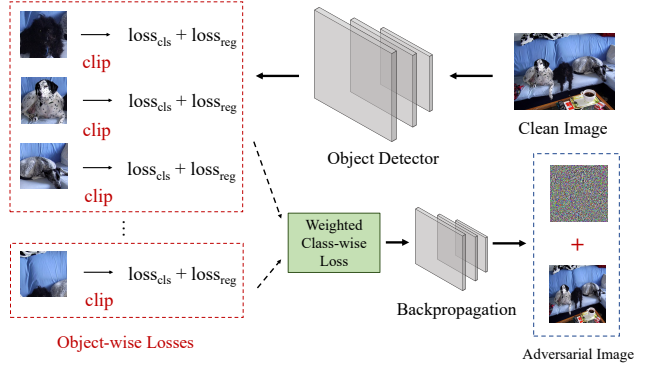


Figure 2. The framework of generating class-wise adversarial examples. In the process of class-wise adversarial generation, we first separate task-oriented losses into object-wise losses and clip each classification and regression loss of an object to force the values of them in the same scale. Then, we generate perturbations from the weighted class-wise loss. Finally, we add the class-wise adversarial perturbations into clean images to generate the class-wise adversarial images.

els for all the object classes as compared with the previous defense methods. In addition, due to the high computational cost of vanilla adversarial training, we also adopt the recent developed fast adversarial training methods [20] into the proposed approach to accelerate the training speed to suffice the practical needs of real-world applications. With extensive experiments on the challenging PASCAL VOC [2] and MS-COCO [12] datasets, the evaluation results demonstrate that the proposed defense methods can effectively enhance the robustness of the object detection models. We summarize the main contributions of our work as follows:

- We provide a systematic analysis and design several efficient and effective adversarial training algorithms for object detection, especially for the situations when there are multiple objects from different classes appearing in a given image. The proposed approaches can craft adversarial examples which can more evenly attack all the objects occurred in an image than previous methods and help improve the adversarial robustness of the trained model with adversarial training.
- We build the connection between the universal adversarial perturbation in the context of image classification and the object detection.

2. Related Works

Due to a large amount of related works in the literature, we briefly review recent relevant works as follows.

2.1. Robustness of the Classification Models

For the adversarial attacks to the deep learning-based models, Szegedy et al. [25] first presented a method to

craft adversarial examples by performing one-step back-propagation given a pretrained classification model, called Fast Gradient Sign Method (FGSM) [5]. FGSM emphasizes the significant impact on the direction of the gradient with respect to the input image. To defend the FGSM attack, they also proposed the adversarial training by using the adversarial images generated by performing the FGSM approach to train the original model. In this way, the worst-case error with the perturbed data would be minimized. From then on, various attacks and defense algorithms have been presented [9], [28] [16], [1], [17], [11], [26] for the classification tasks. Projected Gradient Descent (PGD) is one of the strongest “first-order adversary” attacks as shown in [14] which repeats the optimization of the aforementioned saddle point formulation several times to generate adversarial examples. In addition, training using this min-max formulation makes the learned model more robust toward adversarial attacks.

Unlike “per-instance attack”, recently Seyed-Mohsen Moosavi-Dezfooli et al. [15] proposed universal adversarial perturbations and Shafahi et al. [28] proposed universal adversarial training. For these works, instead of finding per-instance perturbations for each input image, they reformulate the original optimization problem to craft a universal attack to multiple instances at a time across different classes as shown in equation (2). Similarly, based on universal adversarial perturbation, they also proposed corresponding universal adversarial training schemes as shown in equation (3).

$$\max_{\delta} \mathcal{L}(\theta, \delta) = \frac{1}{N} \sum_{i=1}^N l(x_i + \delta, \{y_i\}, \theta) \text{ s.t. } \|\delta\|_p \leq \epsilon \quad (2)$$

$$\min_{\theta} \max_{\|\delta\|_p \leq \epsilon} \mathcal{L}(\theta, \delta) = \frac{1}{N} \sum_{i=1}^N l(x_i + \delta, \{y_i\}, \theta) \quad (3)$$

where $l(\cdot, \cdot, \theta)$ represents the loss function of the training model, δ is the adversarial perturbation, and $\|\delta\|_p \leq \epsilon$ denotes the ℓ_p -norm to prevent δ from growing too large. Moreover, they use a “clipped” version of the loss function,

$$l(x_i + \delta, \{y_i\}, \theta) = \min \{l(x_i + \delta, \{y_i\}, \theta), \beta\} \quad (4)$$

For the universal adversarial perturbation as shown in equation (4), Shafahi et al. constrained the loss value at most β through the clipping operations to prevent the classification loss of any single image from dominating the overall loss of multiple images as shown in equation (2).

2.2. Adversarial Attack for Object Detection

Unlike image classification problem, object detection is the task to detect the occurrence of semantic objects in an image. There are two main categories: one-stage [18, 13] and two-stage object detectors [4, 3, 6, 19]. The detection procedure of a two-stage object detector consists of two steps: (1) the region proposal step and (2) the classification

and localization step. In this paper, we focus on one-stage detectors which simultaneously predict the bounding boxes and classify the anchors in a single inference, and the running speed is much faster than the two-stage ones in solving real-world problems.

Recently, many adversarial attacks are developed for object detection models, and most of them focus on attacking the two-stage detectors. The first attack algorithm is DAG proposed by Xie et al. [29], which specifies the adversarial labels and use back-propagation to iteratively mislead the predictions of the object detectors. Then, Li et al. [10] designed RAP algorithm which combines the label loss and shape loss to yield the adversarial perturbation and optimizes the objective function with an iterative gradient based method. Wei et al. [27] claimed these attack methods focus on attacking the object proposal-based detector have two limitations including weak transferability and high computation cost. Therefore, they proposed UEA to generate adversarial examples using Generative Adversarial Network (GAN) framework and combine with high-level class loss and low-level feature loss.

Although many adversarial attacks have been developed for object detection during the past few years, the defense methods for object detection are rare. Zhang et al. [30] proposed an adversarial training-based algorithm to enhance the robustness of the one-stage detector. Their algorithm decomposes the adversarial training into two task-oriented domains: S_{cls} for the classification branch and S_{reg} for the regression branch of the object detection loss:

$$S_{cls} \triangleq \left\{ x'_{cls} \mid \arg \max_{x'_{cls} \in S_x} l_{cls}(x'_{cls}, \{y_k\}, \theta) \right\} \quad (5)$$

$$S_{reg} \triangleq \left\{ x'_{reg} \mid \arg \max_{x'_{reg} \in S_x} l_{reg}(x'_{reg}, \{b_k\}, \theta) \right\} \quad (6)$$

where x'_{cls} and x'_{reg} represent the adversarial examples generated from each task, S_x is defined as $S_x = \{z \mid z \in \mathcal{B}(x, \epsilon) \cap [0, 255]^n\}$ and $\mathcal{B}(x, \epsilon) = \{z \mid \|z - x\|_{\infty} \leq \epsilon\}$ denotes the ℓ_{∞} -ball with the center as the clean image x and the radius is the perturbation budget ϵ .

Thus, they presented an adversarial training approach according to the task-oriented domain constraint $S_{cls} \cup S_{reg}$ that generates adversarial examples respectively from the object classification and bounding box regression tasks and selects the one which maximizes the overall object detection loss as shown in equation (7).

$$\min_{\theta} \left[\max_{x' \in S_{cls} \cup S_{reg}} l(x', \{y_k, b_k\}, \theta) \right] \quad (7)$$

where x' is the adversarial example $x + \delta$ generated from the clean image x . Since the classification and regression losses are considered independently, the generated adversarial examples may not be able to effectively attack both branches of the object detector.

3. Methodology

To generate the adversarial examples which can effectively and evenly fool the object detector to change the detection results of all the occurred objects in an image, we develop a novel class-aware robust adversarial training for the object detection task. The proposed approach considers *heterogeneous* (classification and regression tasks), *multiple* (multiple objects), and *balanced* (multiple classes) class losses to generate adversarial examples for a robust object detector. The overview of the proposed framework for generating class-wise adversarial examples is shown in Figure 2 and the details of the proposed approach are described as follows:

3.1. Multi-task Adversarial Training for Object Detection

To address the issues of adversarial training using the overall object detection loss, we delve into the details of the object detection loss. Different from the classification task which only contain a single loss (e.g., cross entropy loss for the classification task) to predict the results, the loss of the one-stage detector consists of two different kinds of losses: (1) the classification loss for predicting the category scores and (2) the regression loss (e.g., smooth L_1 loss) for predicting the box offsets from the input images to detect objects. Therefore, we have *heterogeneous* sources of losses from the classification task and regression task for the generation of adversarial examples and adversarial training. We can define the following optimization problem for building the multi-task adversarial training for object detection.

$$\min_{\theta} \max_{\|\delta\|_p \leq \epsilon} \mathcal{L}(\theta, \delta) = \hat{l}_{cls}(x + \delta, \{y\}, \theta) + \hat{l}_{reg}(x + \delta, \{b\}, \theta) \quad (8)$$

where $\hat{l}_{t \in \{cls, reg\}}(\cdot, \{y, b\}, \theta)$ represents the confidence loss and localization loss used for the one-stage detectors respectively. Furthermore, the naive loss function as shown in equation (1) suffers from a significant impact that each task-oriented loss is unbounded, and one of the task-oriented losses can be extremely large during the adversarial training process. In the worst case, the value of the loss might go to infinity and dominate the overall object detection loss. To address this issue, we propose a “clipped” version for each task-oriented loss function,

$$\hat{l}_{t \in \{cls, reg\}}(x + \delta, \{y\}, \theta) = \min \left\{ \hat{l}_t(x + \delta, \{y\}, \theta), \beta_t \right\} \quad (9)$$

As we have shown in equation (9), this method can not only avoid the mutual interference between each task but also prevent any task-oriented loss from dominating the overall objective function for object detection by regularizing each task-oriented loss values at most β_t . In the experimental result section, we will perform an ablation studies to show the effect of clipping the loss with different thresholds, β_t .

With this objective function, it will search an adversarial perturbation which can jointly and effectively maximizes both task-oriented losses for all the objects occurred in an image.

3.2. Object-wise Adversarial Training for Object Detection

Besides the task-oriented losses, we propose the second objective function that further delves into the object detection loss in the *multiple* object aspect by considering the scenarios that there usually exist multiple objects in an image for the object detection task. We thus propose the object-wise adversarial training for object detection and formulate this problem as a min-max optimization problem as follows:

$$\min_{\theta} \max_{\|\delta\|_p \leq \epsilon} \mathcal{L}(\theta, \delta) = \sum_{i=1}^{N_o} \hat{l}_{cls}^o(O_i + \delta, \{y_i\}, \theta) + \sum_{i=1}^{N_o} \hat{l}_{reg}^o(O_i + \delta, \{b_i\}, \theta) \quad (10)$$

As we discuss in the multi-task adversarial training, the clipped version of the loss function would prevent the single task-oriented loss function to dominate the overall training loss. Since each task-oriented loss for a specific object (i.e., the loss of a specific object goes to infinity.) might dominate the overall loss value during the generation process of adversarial examples, we similarly propose a “clipped” version for the proposed object-wise loss function as shown in equation (11) which can bound the task-oriented losses of each object at most β_o . We can then generate the adversarial perturbation for the multiple objects in an image using a similar way as the universal adversarial attack proposed in [28] for the attacks of multiple images in the classification setting.

$$\hat{l}_{t \in \{cls, reg\}}^o(O_i + \delta, \{y_i\}, \theta) = \min \left\{ \hat{l}_t(O_i + \delta, \{b_i\}, \theta), \beta_o \right\} \quad (11)$$

3.3. Class-wise Adversarial Training for Object Detection

Furthermore, for the object detection task, it usually not only contains multiple objects but various classes for a given image. To prevent the loss of a specific object class from dominating the overall object detection loss (i.e., in a given image, there are more objects of a specific class than other classes.), we propose the third objective function as shown in equation (12). Instead of normalizing the total loss with the number of objects, the proposed approach decomposes the total loss into class-wise task-oriented losses and normalizes each class loss using the number of objects for the corresponding class in a given image. The adversarial training based on the proposed class weighted loss can effectively balances the influence of each class during ad-

versarial training.

$$\mathcal{L}_{c'} = \frac{1}{C} \sum_{i=1}^C \frac{1}{n_i} \sum_{j=1}^{n_i} \hat{l}_{cls}^o(O_j, \{y_j\}, \theta) + \hat{l}_{reg}^o(O_j, \{b_j\}, \theta) \quad (12)$$

where C is the number of classes in one image, n_c is the number of the matched default boxes in the class c .

3.4. Fast Adversarial Training for Object Detection

Since the proposed defense methods focus on the complex object detection models, it requires much more computational resources to train a model than those of an image classification model, and the efficiency of the adversarial training is a critical point to train a robust detector within a reasonable amount of time and computational resources for the real world applications. However, the high computational cost of the iterative gradient back-propagation of the PGD-based adversarial training makes it less practical. Shafahi et al. [20] recently proposed a fast adversarial training algorithm which recycles the gradient information to reduce the cost of generating adversarial examples when updating the model parameters. With the fast adversarial training, the training process could be 7 to 30 times faster than the original adversarial training. We thus adopt it into the proposed algorithm. Finally, the final version of the proposed algorithm takes the *heterogeneous* tasks, *multiple* objects, *balanced* class losses and fast training into consideration, and the details of the proposed fast class-wise adversarial training for object detection are summarized in Algorithm 1. Due to limited space, we also refer the readers to Appendix for more running time analysis for the fast adversarial training.

Algorithm 1 Fast Class-wise Adversarial Training

Require: dataset D , training epoch N_{ep} , perturbation bound ϵ , learning rate γ

- 1: **for** epoch = 1, ..., N_{ep}/m **do**
- 2: **for** minibatch $B \sim D$ **do**
- 3: **for** iter = 1 to m **do**
- 4: Compute gradient of loss with respect to δ
- 5: $d_\delta \leftarrow \mathbb{E}_{x \in B} [\nabla_\delta \mathcal{L}_{c'}(\theta, x + \delta)]$
- 6: Update θ with momentum stochastic gradient
- 7: $g_\theta \leftarrow \mu g_\theta - \mathbb{E}_{x \in B} [\nabla_\theta \mathcal{L}(\theta, x + \delta)]$
- 8: $\theta \leftarrow \theta + \gamma g_\theta$
- 9: Update perturbation δ with gradient
- 10: $\delta \leftarrow \delta + \epsilon \text{sign}(d_\delta)$
- 11: Project δ to ℓ_p -ball
- 12: **end for**
- 13: **end for**
- 14: **end for**

4. Experimental Results

In this section, we first describe the experimental settings and then show the evaluation results of the proposed adversarial defense approaches for the object detectors on the challenging PASCAL VOC [2] and MS-COCO [12] datasets in the following sections.

4.1. Datasets and Evaluation Settings

For the PASCAL VOC dataset, we adopt the standard “07+12” protocol, which contains a total of 16,551 images, 40,058 objects, and 20 classes for training. On the testing phase, we use the test set of the PASCAL VOC 2007 dataset with a total of 4,952 testing images. To evaluate the performance of the object detector after adversarial attacks, we compute the average precision (AP) for the category of interest and the “mean average precision” (mAP) for the overall performance. For the MS-COCO dataset, we train the model using its training set in 2017 with a total of 118,287 images, and the number of object categories is 80. For testing, we evaluate the results using its validation set in 2017 with a total of 5,000 images. The mAP with IoU threshold 0.5 is used for evaluating the robustness of a detector.

In addition, we also introduce the algorithms for comparison in all the experiments as follows:

- **STD**: the object detector trained with natural training using clean images.
- **CLS**¹: the model trained using A_{cls} for PGD-based adversarial training.
- **REG**¹: the model trained using A_{reg} for PGD-based adversarial training.
- **MTD**¹/**MTD-fast**: the model that we trained with our own implementation with normal/fast adversarial training in [30] where we denote the generated adversarial examples as multi-task domain attack (MDA).
- **TOAT**: the model trained with the proposed task-oriented PGD-based adversarial training.
- **OWAT**: the model trained with the proposed object-wise PGD-based adversarial training where we also denote the generated adversarial examples as the object-wise attack (OWA).
- **CWAT**: the model trained with the proposed class-wise PGD-based adversarial training where we also denote

¹Since the official implementation is not available, we re-implement it with the same setting as presented in [30]. In addition, since the paper did not describe the parameter, the number of steps for the PGD-based adversarial training, we can only choose the one closest to their results as the foundation of performance comparison, and we also present the original result of CLS, REG, and MTD.




Figure 3. The robustness of each model under attacks with $\epsilon = 8$ from different number of steps in PASCAL VOC 2007 test set.

the generated adversarial examples as the class-wise attack (CWA).

4.2. Implementation Details

We conduct experiments using the one-stage detectors, SSD, with the VGG-16 network backbone as the main evaluation test bed where we use a modified version of the VGG-16 with batch normalization layers. All the models used in the experiments are fine-tuned from the pretrained SSD model using the training set of the corresponding object detection benchmarks and the SGD optimizer with an initial learning rate, 10^{-2} , momentum, 0.9, and weight decay 0.0005 with the multi-box loss. The learning rates are decayed at 16th and 20th epochs respectively with the decay factor equal to 0.1. The resolution of the resized input image is 300×300 . The range of the pixel values is between $[0, 255]$ and then shifted according to the mean of pixel intensities of the whole dataset. For adversarial training, we use the budget $\epsilon = 8$ to generate the adversarial examples as the inputs. To be more specific, we denote A_{cls} as the classification-task adversarial examples generated only considering the overall classification loss and A_{reg} as the regression-task adversarial examples generated only considering the overall regression loss. For our fast PGD-based adversarial training, we set $m = 4$. In addition, the SSD and its adversarial robust version are also trained with the online hard example mining strategy (OHEM) [22] to sample hard negative samples. We set $\beta_t = \beta_o = 6$ for all the experiments in this paper.

4.3. Evaluation Results on Pascal VOC and MS-COCO

In this subsection, we show the evaluation results on both the Pascal VOC and MS-COCO datasets, where the MS-COCO dataset is more close to the real-world object detection scenarios and more challenging for testing on the robustness of object detectors. The results of different models under the PGD-10 attack with attack budget $\epsilon = 8$, and other attacks, including FGSM, CWA and DAG², are summarized in Table 1 and 2. We can find that

²We implement DAG algorithm into our one-stage models, the results are similar with the UEA experiments.

attack	clean	FGSM		PGD-10		CWA	DAG
		A_{cls}	A_{reg}	A_{cls}	A_{reg}		
STD	0.752	0.162	0.25	0.012	0.043	0.006	0.291
CLS	0.467	0.309	0.343	0.236	0.334	-	-
REG	0.519	0.263	0.344	0.146	0.279	-	-
MTD-fast	0.466	0.311	0.418	0.221	0.351	0.182	0.486
TOAT-6	0.430	0.300	0.397	0.218	0.334	0.197	0.466
OWAT	0.518	0.327	0.434	0.229	0.372	0.203	0.504
CWAT	0.513	0.325	0.433	0.224	0.367	0.199	0.503

Table 1. The evaluation results of various adversarial trained SSD model with the VGG16-BN backbone network under FGSM PGD-10 attacks with $\epsilon = 8$, CWA, and DAG in PASCAL VOC 2007 test set.

attack	clean	FGSM		PGD-10		CWA
		A_{cls}	A_{reg}	A_{cls}	A_{reg}	
STD	0.451	0.133	0.167	0.030	0.029	0.003
MTD ¹	0.190	0.127	0.146	0.110	0.135	0.082
MTD-fast	0.242	0.167	0.182	0.130	0.134	0.077
TOAT-6	0.182	0.120	0.148	0.098	0.123	0.074
OWAT	0.211	0.129	0.169	0.100	0.140	0.074
CWAT	0.237	0.168	0.189	0.142	0.155	0.092


Table 2. The adversarial robustness of each model using SSD VGG16-BN model under FGSM, PGD-10, and CWA attacks with $\epsilon = 8$ in the MS-COCO test set.

for the Pascal VOC dataset, the proposed OWAT achieves the best performance than other compared methods while CWAT achieves comparable performance as compared with OWAT. For the MS-COCO dataset, CWAT achieves the highest performance under the PGD-10 adversarial attack over MTD-fast, TOAT, and OWAT. As shown in Figure 4, most of the images in the Pascal VOC contains much fewer objects than the MS-COCO, especially in different classes. Note that CWAT is less effective when the number of classes in a single image is few. That is why CWAT just achieves comparable performance to OWAT. The evaluations result of the MS-COCO dataset confirms the CWAT could effectively balance class loss to prevent the loss of a specific class from dominating the overall object detection loss, especially when we focus on comparing the results of CWAT and OWAT.


4.4. Ablation Study

4.4.1 Attack under Different Number of PGD Steps and Different Budgets

To evaluate the performance of the proposed adversarial training for object detection and compare with previous methods, we first attack the models using the adversarial examples generated with different number of PGD steps. As shown in Figure 3, the proposed OWAT and CWAT both can enhance the robustness for these settings. With the proposed CWAT, the performance can be significantly enhanced as compared with our implemented MTD-fast where



(a) PASCAL VOC



(b) MS-COCO

Figure 4. The accumulation of number of objects and number of categories contained in an image, and the total results of the number of images toward different compositions.




Figure 5. The detection results of the class-wise adversarial attack with PGD-10 in different ϵ , using the clean SSD as the targeted model. White label, yellow label, magenta label, blue label, green label, and black label represent classes of the dining table, chair, potted plant, sofa, tv-monitor, and, respectively. (The more detailed qualitative and quantitative results can be found in Appendix.)

MTD [30] is the recent state-of-the-art adversarial training method for object detection. In addition, by taking both training time and the training settings of other related works into consideration, we choose PGD-10 to generate the adversarial examples for training. Moreover, we also evaluate the robustness of models under different budgets as shown in Figure 5. We also evaluate each model under the adversarial attacks with different budgets, and we summarize the results in Appendix.

4.5. The Impact of Different Thresholds β_{cls} and β_{reg} in Multi-task Adversarial Training

To analyze the effect of the “clipping” parameters of β_{cls} and β_{reg} , we further explore the proposed TOAT in different clipping parameters as follows.

- **TOAT-inf**: the model trained using TOAT and with the values of β_{cls} and β_{reg} set as ∞ .

- **TOAT-i**: the model trained using TOAT and with the values of β_{cls} and β_{reg} as i .
- **TOAT-k-l**: the model trained using TOAT and with the values of β_{cls} as k and β_{reg} as l .

As shown in Table 5, we can find that β_{cls} and β_{reg} have important impacts on the adversarial robustness. As a result, the “clipping” parameters can prevent one of the task-oriented losses from dominating by the other. To choose the proper clipping thresholds, we first run two experiments with same β_{cls} and β_{reg} called TOAT-i. For example, we set $\beta_{cls} = \beta_{reg} = 2$ and $\beta_{cls} = \beta_{reg} = 4$ to analyze the impact on different parameters. We can explore that TOAT-2 has better performance defended against l_{reg} attacks. On the other hand, TOAT-4 has better performance defended against l_{cls} attacks. Therefore, if β_{cls} and β_{reg} are in different settings and and balance both task-oriented losses, we can optimize the adversarial training and get the better robustness. Therefore, we test TOAT with different β_{cls} and β_{reg} called TOAT-k-l. After a simple testing process, we choose the clipping parameters $\beta_{cls} = \beta_{reg} = 6$ as the final set and implement it into the OWAT and CWAT models.

4.6. The Impact of Task-oriented Attack and Class-wise Attack for Object Detection

As described in the methodology section, not all of the objects in an image would be attacked successfully. To analyze the potential influences between each object in an image, we explore this impact by delving into the object losses with adversarial attack in each class and by comparing the clean SSD model with CWAT.

	Clean SSD			OWAT SSD			CWAT SSD		
	MDA	OWA	CWA	MDA	OWA	CWA	MDA	OWA	CWA
aeroplane	0.010	0.019	0.002	0.383	0.370	0.365	0.354	0.373	0.353
bicycle	0.024	0.092	0.001	0.354	0.365	0.357	0.334	0.340	0.338
bird	0.003	0.003	0.001	0.118	0.117	0.103	0.106	0.110	0.111
boat	0.006	0.001	0.000	0.186	0.174	0.126	0.165	0.161	0.116
bottle	0.037	0.019	0.001	0.101	0.101	0.099	0.101	0.102	0.100
bus	0.023	0.011	0.001	0.292	0.316	0.278	0.298	0.312	0.266
car	0.020	0.010	0.092	0.409	0.423	0.425	0.411	0.429	0.435
cat	0.000	0.001	0.000	0.148	0.162	0.100	0.146	0.157	0.087
chair	0.005	0.004	0.001	0.095	0.095	0.109	0.109	0.107	0.109
cow	0.000	0.002	0.000	0.054	0.050	0.058	0.089	0.069	0.033
diningtable	0.005	0.005	0.000	0.317	0.317	0.220	0.288	0.290	0.207
dog	0.003	0.000	0.000	0.080	0.083	0.071	0.116	0.125	0.093
horse	0.007	0.003	0.002	0.369	0.373	0.277	0.336	0.329	0.284
motorbike	0.002	0.003	0.002	0.317	0.328	0.277	0.301	0.317	0.265
person	0.035	0.012	0.016	0.317	0.317	0.311	0.314	0.319	0.316
pottedplant	0.007	0.091	0.000	0.099	0.097	0.095	0.096	0.096	0.094
sheep	0.001	0.012	0.001	0.144	0.150	0.161	0.173	0.156	0.149
sofa	0.001	0.000	0.000	0.155	0.174	0.134	0.156	0.160	0.142
train	0.018	0.002	0.001	0.315	0.321	0.259	0.251	0.262	0.248
tvmonitor	0.007	0.004	0.002	0.249	0.251	0.241	0.253	0.264	0.235
mAP	0.011	0.015	0.006	0.225	0.229	0.203	0.220	0.224	0.199

Table 3. The average precision of each category under different PGD-10 attacks including MDA, OWA, and OWA with different robust SSD models in PASCAL VOC 2007 test set.

The Robustness of Different SSD Models on Each Category: To verify the effectiveness of the proposed approaches and to improve the adversarial robustness for each object class evenly, we show the per-class performances after various adversarial attacks. As shown in Table 3, the proposed approach significantly improves the adversarial robustness for each class, which also proves the effectiveness of our approach.

The Effectiveness of Different Attack Methods on Each Category: To delve into the effectiveness of different proposed attack methods, we can also figure out by the per-class performances after defense these adversarial attacks. As shown in Table 3, the MDA proposed by [30] is similar to the proposed OWA in this paper. Additionally, the proposed CWA is much more powerful than others, which also proves our concept mentioned in the introduction section that the class-wise adversarial attack can attack each class in a balanced way. However, there are few categories OWAT and CWAT will fail such as bicycle and car. In our observation, the categories which will cause the attack fail almost present as a large object in the images. Due to the SSD limitation of the large scale detection, it may cause the attack fail toward the large objects.

4.7. Evaluation using Different Network Architecture

We also evaluate the SSD models using different backbone networks. In addition to the original VGG-16BN backbone, we also use the SSD model with the ResNet-50 as a backbone from the implementation of ScratchDet [31]. As shown in Table 4, the proposed method can improve the robustness of the SSD models significantly by 20% to 30% mAP across different backbone networks, demonstrating that the proposed CWAT can consistently improve the

attack	clean	FGSM		PGD-10		CWA	DAG	
		A_{cls}	A_{reg}	A_{cls}	A_{reg}			
VGG16	STD	0.752	0.162	0.250	0.012	0.043	0.006	0.291
	CWAT	0.513	0.324	0.433	0.222	0.366	0.199	0.503
ResNet50	STD	0.806	0.344	0.420	0.070	0.062	0.016	0.795
	CWAT	0.483	0.318	0.401	0.244	0.340	0.250	0.482

Table 4. The robustness of different one-stage models including SSD with different backbone networks under FGSM, PGD-10, and CWA attacks with $\epsilon = 8$ in the PASCAL VOC 2007 test set.

attack	clean	FGSM		PGD-10		CWA	DAG
		A_{cls}	A_{reg}	A_{cls}	A_{reg}		
TOAT-inf	0.469	0.266	0.374	0.173	0.310	0.165	0.456
TOAT-2	0.535	0.302	0.408	0.191	0.297	0.132	0.502
TOAT-4	0.477	0.288	0.396	0.190	0.327	0.183	0.463
TOAT-6	0.430	0.300	0.397	0.218	0.334	0.197	0.466
TOAT-8	0.343	0.289	0.389	0.204	0.331	0.189	0.457
TOAT-2-4	0.512	0.304	0.388	0.211	0.292	0.131	0.487
TOAT-4-2	0.479	0.296	0.404	0.203	0.337	0.178	0.472
TOAT-2-6	0.529	0.325	0.403	0.226	0.302	0.146	0.503
TOAT-6-2	0.467	0.289	0.382	0.204	0.316	0.175	0.452

Table 5. The evaluation results of different clipping hyperparameters of the proposed TOAT for the classification and regression losses for the adversarially trained SSD model with the VGG16-BN backbone network under FGSM PGD-10 attacks with $\epsilon = 8$, CWA, and DAG in PASCAL VOC 2007 test set.

adversarial robustness of the object detectors with different backbone networks.

5. Conclusion

In this work, we present several robust adversarial training for object detection. For a given image, the proposed approach can generate an effective universal adversarial perturbation to simultaneously attack all the occurred objects in the image through jointly maximizing the respective loss for each object. Additionally, the proposed class-wise adversarial training for object detection can not only balances the influence of each class but also effectively and evenly improves adversarial robustness of trained models for all the object classes as compared with the previous defense methods. Meanwhile, with the recent development of fast adversarial training, we provide a fast version of the proposed algorithm, which can be trained faster than the traditional adversarial training while keeping performance comparable. With extensive experiments on the challenging PASCAL-VOC and MS-COCO datasets, the evaluation results demonstrate that the proposed defense methods can effectively enhance the robustness of the object detection models.

Acknowledgement This work is supported by Ministry of Science and Technology (MOST), Taiwan, under Grants No. 108-2218-E-001-004-MY2, 109-2221-E-001-020-, and 109-2221-E-001-016-.

References

- [1] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57, 2017.
- [2] Mark Everingham, S. M. Ali Eslami, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision (IJCV)*, 111:98–136, 2014.
- [3] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision (ECCV)*, pages 1440–1448, 2015.
- [4] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 580–587, 2014.
- [5] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2015.
- [6] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision (ICCV)*, pages 2961–2969, 2017.
- [7] Yi Yao Huang and William Yang Wang. Deep residual learning for weakly-supervised relation extraction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1803–1807, 2017.
- [8] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems (NeurIPS)*, pages 1097–1105, 2012.
- [9] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*, 2016.
- [10] Yuezun Li, Daniel Tian, Ming-Ching Chang, Xiao Bian, and Siwei Lyu. Robust adversarial perturbation on deep proposal-based models. In *British Machine Vision Conference (BMVC)*, page 231, 2018.
- [11] Fangzhou Liao, Ming Liang, Yinpeng Dong, Tianyu Pang, Xiaolin Hu, and Jun Zhu. Defense against adversarial attacks using high-level representation guided denoiser. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1778–1787, 2018.
- [12] 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 *European conference on computer vision (ECCV)*, pages 740–755, 2014.
- [13] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. SSD: Single shot multibox detector. In *European conference on computer vision (ECCV)*, pages 21–37, 2016.
- [14] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations (ICLR)*, 2018.
- [15] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. Universal adversarial perturbations. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 1765–1773, 2017.
- [16] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 2574–2582, 2016.
- [17] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *IEEE Symposium on Security and Privacy (SP)*, pages 582–597. IEEE, 2016.
- [18] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- [19] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 91–99, 2015.
- [20] Ali Shafahi, Mahyar Najibi, Mohammad Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3358–3369, 2019.
- [21] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K Reiter. Adversarial generative nets: Neural network attacks on state-of-the-art face recognition. *arXiv preprint arXiv:1801.00349*, 2017.
- [22] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 761–769, 2016.
- [23] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015.
- [24] Chawin Sitawarin, Arjun Nitin Bhagoji, Arsalan Mosenia, Mung Chiang, and Prateek Mittal. Darts: Deceiving autonomous cars with toxic signs. *arXiv preprint arXiv:1802.06430*, 2018.
- [25] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*, 2014.
- [26] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian J. Goodfellow, Dan Boneh, and Patrick D. McDaniel. Ensemble adversarial training: Attacks and defenses. In *International Conference on Learning Representations (ICLR)*, 2018.
- [27] Xingxing Wei, Siyuan Liang, Ning Chen, and Xiaochun Cao. Transferable adversarial attacks for image and video object detection. In *Proceedings of the 28th International Joint*

- Conference on Artificial Intelligence (IJCAI)*, pages 954–960, 2019.
- [28] Rey Wiyatno and Anqi Xu. Maximal jacobian-based saliency map attack. *arXiv preprint arXiv:1808.07945*, 2018.
 - [29] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Yuyin Zhou, Lingxi Xie, and Alan Yuille. Adversarial examples for semantic segmentation and object detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1369–1378, 2017.
 - [30] Haichao Zhang and Jianyu Wang. Towards adversarially robust object detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 421–430, 2019.
 - [31] Rui Zhu, Shifeng Zhang, Xiaobo Wang, Longyin Wen, Hailin Shi, Liefeng Bo, and Tao Mei. Scratchdet: Training single-shot object detectors from scratch. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 2268–2277, 2019.