arXiv:2112.14232v1 [cs.LG] 28 Dec 2021

Constrained Gradient Descent: A Powerful and Principled Evasion Attack Against Neural Networks

Weiran Lin
Carnegie Mellon University
weiranl@andrew.cmu.edu

Keane Lucas
Carnegie Mellon University
kjlucas@andrew.cmu.edu

Lujo Bauer
Carnegie Mellon University
lbauer@cmu.edu

Michael K. Reiter
Duke University
michael.reiter@duke.edu

Mahmood Sharif
Tel Aviv University
mahmoods@cs.tau.ac.il

Abstract

Minimal adversarial perturbations added to inputs have been shown to be effective at fooling deep neural networks. In this paper, we introduce several innovations that make white-box targeted attacks follow the intuition of the attacker's goal: to trick the model to assign a higher probability to the target class than to any other, while staying within a specified distance from the original input. First, we propose a new loss function that explicitly captures the goal of targeted attacks, in particular, by using the logits of all classes instead of just a subset, as is common. We show that Auto-PGD with this loss function finds more adversarial examples than it does with other commonly used loss functions. Second, we propose a new attack method that uses a further developed version of our loss function capturing both the misclassification objective and the L_∞ distance limit ϵ. This new attack method is relatively 1.5–4.2% more successful on the CIFAR10 dataset and relatively 8.2–14.9% more successful on the ImageNet dataset, than the next best state-of-the-art attack. We confirm using statistical tests that our attack outperforms state-of-the-art attacks on different datasets and values of ϵ and against different defenses.

1 Introduction

With the prevalence of machine learning, adversarial-machine-learning techniques that slightly manipulate the inputs of a machine-learning model to influence its functionality have also been developed [14]. One type of attack on classification models is the *evasion* attack, which applies a small perturbation, within a distance limit, to a classifier input to cause the classifier to misclassify. The L_∞ *distance* is commonly used for specifying distance limits via ϵ *boundaries*, which define the maximum change acceptable in each element of the input. It is also a common setup that attackers have access to all weights of the model, known as the *white-box* scenario, and the attackers aim to force specific misclassification, performing *targeted* attacks. Previous work proposes an attack method that iteratively moves in the direction pointed by the gradients of a loss function, which is not necessarily the same loss function used by the model, and automatically truncates the attack in each iteration to stay within the distance limits. Researchers have shown that such attacks are effective against state-of-the-art neural networks [35, 50].

It is necessary to understand how good attacks work and look for stronger ones. Attacks serve as one of the metrics to evaluate goodness of the models [54]. These attacks could help humans to understand how vulnerable the models are and build better defenses, systems that are robust to real world scenarios where small perturbations might change the predictions of models. Targeted attacks, adversarial perturbations that enforce specific misclassification, are especially interesting as they correspond to certain attacker goals and potentially cause higher damage compared to attacks that only avoid the correct classification. Meanwhile, white-box attacks, attacks where the adversary has full access to the internal weights of the models, are also worth paying attention to. Some real world systems, such as open-source models, are already in white-box settings. The white-box scenario defines the worst case scenario in reality and could help to defend both strong and weak attacks. It is common among previous works to study white-box targeted attacks [14, 19, 33]. This paper reveals the process of finding stronger white-box targeted attacks within the L_∞ distance limit.

Croce et al. show that well-tuned parameters and carefully designed loss functions can boost the performance of attacks [14]. For example, varying the loss function of Auto Projected Gradient Descent (Auto-PGD), a state-of-the-art attack, between the cross-entropy, the Carlini and Wagner (CW), and the Difference of Logits Ratio (DLR) losses has a substantial impact on the attack's performance [14]. Guided by this observation to improve the performance of targeted attacks, we define a new loss function, the *Minimal Difference loss (MD loss)*, that best aligns with the attacks' goal: the MD loss aims to (mis)lead the model to assign higher confidence to the target class than to *any other class*, even if by just a tiny amount. Due to best capturing the adversary's

goal, we empirically observe that Auto-PGD with the MD loss finds on average 0.5%–12.3% more adversarial examples compared to Auto-PGD with other loss functions, depending on the specific dataset and model.

Although the MD loss significantly improves Auto-PGD's performance, we still notice certain limitations that hinder the attack's effectiveness. In particular, consistently with the PGD family of attacks, Auto-PGD uses projection at the end of each iteration to satisfy the norm constraints, eliminating changes that fall outside the limits of a predefined L_p ball. However, projecting adversarial examples back into the L_p ball may contradict the attacker's misclassification objective, thus harming the attack's success. Moreover, because projection is implemented by clipping in each iteration, an *ad hoc* operation to eliminate any changes made outside the L_p ball, balancing the attack's two objectives (i.e., misclassification and not exceeding the L_p ball) remains a challenge for Auto-PGD. To address this shortcoming, we propose the *Constrained Gradient Descent (CGD)* attack. CGD enables the attacker to balance the two attack objectives—it allows adversarial examples to lie outside L_p ball during the attack process and models the violation of the norm constraints as part of its loss function to gradually learn to stay within the L_p ball while achieving misclassification. We find that, within the same L_∞ distance limit, CGD finds on average 0.3%–1.2% more adversarial examples compared to Auto-PGD using the MD loss. Furthermore, we verify that CGD outperforms Auto-PGD with three previously established loss functions on the CIFAR10 and ImageNet datasets, with statistical significance.

As an example of using CGD as a framework to find other types of attacks, we also define a variant of CGD as white-box untargeted attacks. Similar to its targeted attack counterpart, this variant also allows adversarial examples to lie outside L_p ball during the attack process and models the violation of the norm constraints as part of its loss function to gradually learn to stay within the L_p ball while achieving misclassification. This variant outperforms the previous best attack by 0.25%–2% on the CIFAR10 dataset, on average.

To summarize, in this work, our contributions are:

- We improve an established targeted attack method by replacing its loss function with a new loss function that better captures the goal of targeted attacks.
- We invent a new attack that learns to stay within the L_∞ distance limit rather than using simple clipping, and we show that this is a stronger attack compared to previously established ones.

The remainder of the paper proceeds as follows. Next we present background and related work (Sec. 2), with a primary focus on targeted white-box attacks and defenses against them. Subsequently, we introduce our new MD loss (Sec. 3) and the CGD attack (Sec. 4). We then explain the setup of our experiments to establish a standard for comparing attacks (Sec. 5) and report on the evaluation results (Sec. 6). Finally, we explain how CGD can be used as a framework for attacks, and we illustrate this use by instantiating CGD to define a stronger untargeted attack (Sec. 7).

2 Background

We first describe the constraints on acceptable perturbations resulting from our attack and the adversary's knowledge and goals (Sec. 2.1). We then introduce the evaluation metric we choose for attacks (Sec. 2.2). We describe several state-of-the-art attacks (Sec. 2.3) and give an overview of defenses (Sec. 2.4).

2.1 Threat Models

We briefly review some common attack characteristics regarding constraints and adversarial capabilities to better position our attack in the context of prior work.

Distance metrics In creating an adversarial example, most attacks attempt to minimize the changes to the benign image so that the resulting adversarial example is hopefully imperceptibly different to humans [45, 47, 50]. Hence, they use certain distance metrics and set a maximum distance limit ε. One common distance metric is the L_∞ distance [13, 14, 19, 33, 61], which is the maximum change made across all pixels and channels:

$$L_\infty(x, x') = \max_{i,j,k} \left| x_{i,j,k} - x'_{i,j,k} \right| \tag{1}$$

where x' is the adversarial example, x is the benign image which is usually normalized to the range $[0, 1]$, i and j are pixel coordinates, and k is the index number of channels. Meanwhile, with $x_{i,j,k} \in [0, 1]$, the value range a channel (of a pixel) could convey, $x'_{i,j,k}$ is also bounded by $[0, 1]$, in addition to $[x'_{i,j,k} - \varepsilon, x'_{i,j,k} + \varepsilon]$.

Although all norms have their limitations [45, 47, 52], in this work, we particularly study attacks that use the L_∞ distance as it prevents large perturbations that are focused on a small area and thus can erase objects from images [19]. We also find the most defenses, i.e., adversarially trained models, using L_∞ distance metric (e.g., [9, 12, 14, 16, 23, 44, 55, 57, 58]).

Adversary knowledge The amount of information available to the adversary about the model that is going to be attacked decides the optimal approach for the adversary to take [6, 39]. Attacker access to the model is generally categorized as three types [36, 39]: black-box, where the attacker has no information about the model [2, 10, 38]; grey-box, where the attacker has access to the architecture and training but not the specific instance [59]; and white-box, where the attacker has access to the same instance and its internal weights [13, 19, 27]. We study white-box scenarios where

the attacker has the most privileges and hence can show the effectiveness of different attack methods better.

Adversary goals In a supervised classification problem, given the correct label y of an example x, a machine learning model F with loss function L would try to minimize $L(x,y)$ so that $F(x) = y$. This model F is known as the benignly trained model and x is a benign input example. The adversary could either launch an untargeted attack, avoiding correct classification by maximizing $L(x,y)$, or launch a targeted attack, forcing specific classification to class t by minimizing $L(x,t)$ [39]. A classifier computes logits and generates confidence for each class as:

$$P_i = \frac{e^{Z_i}}{\sum_{j=1}^{K} e^{Z_j}}, i = 1..., K \tag{2}$$

Relatively higher logits yield relatively higher confidence, and the classifier selects the class which has the highest confidence as its prediction. The loss function of the classifier, $L(x,y)$ is the cross-entropy loss (CE loss):

$$L_{CE}(x,y) = -\log P_y = -Z_y + \log(\sum_{j=1}^{K} e^{Z_j}) \tag{3}$$

As an example of $L(x,t)$, previous work found that adversaries can generate targeted adversarial examples by replacing the correct label y with the target class t in the cross-entropy loss, thus defining the loss for the attack as [19]:

$$L_{CE}(x,t) = -\log P_t = -Z_t + \log(\sum_{j=1}^{K} e^{Z_j}) \tag{4}$$

Given a specific L_p distance limit ε, the success rate of an untargeted attack is computed as the percentage of benign input examples x from which the attack finds $x' \in \{X' \mid F(X') \neq y, L_p(X',x) \leq \varepsilon\}$, whereas the success rate of a targeted attack is defined as the percentage of benign input examples x where the attack finds $x' \in \{X' \mid F(X') = t, L_p(X',x) \leq \varepsilon\}$. In this paper, we specifically study targeted attacks as Croce and Hein [14] found that targeted attacks could be converted into untargeted attacks by launching targeted attacks towards all the incorrect classes and, if any of these attacks succeed, count it as one untargeted success. This approach found more adversarial examples than the loss function described in Eqn. 3 in their experiments.

2.2 Evaluation Metrics

Previous works evaluated targeted attacks with different metrics. Some previous work evaluated attacks by the distance between successfully perturbed images and their corresponding original benign images [8]. Other work converted targeted attacks to untargeted attacks, [14], as described in Sec. 2.1. In our experiments, we evaluated targeted attacks by success rate as described in Sec. 2.1.

Successful adversarial examples should also be in the same format as benign examples. Images from the two datasets we use (see Sec. 5.1) are in 8-bit RGB format: every pixel consists of three bytes, three integers $\in [0,255]$ that we normalized to floats $\in [0,1]$, as is typical. The value of every channel (of every pixel) should be a multiple of $1/255$. We picked ε values that are multiples of $1/255$ so that the L_∞ distance limit is in 8-bit RGB format. We noticed that PGD attacks could perturb an adversarial example into a failed one. Hence, in each iteration, we produced an 8-bit RGB format copy of the current adversarial example x'_i, which is in the continuous domain, using the formula:

$$x_{test} = round(x'_i * 255)/255 \tag{5}$$

We used $clip(\cdot)$ to clip images within both $[0,1]$ and $[x-\varepsilon, x+\varepsilon]$, making a projection onto the imperfect L_∞ ball that is also bounded by the range of values of valid images. We then classified the clipped adversarial example $clip(x_{test})$ using the model and compared the output with the target class, If they matched, we stopped perturbing this example and counted the attack as successful. Otherwise, the attack continued normally, with x'_i remaining in the continuous domain.

2.3 Attack Methods

Here we introduce several attacks established by previous works.

FGSM Goodfellow et al. proposed the fast gradient sign method (FGSM) attack, which assumes linearity of the classifier and uses the sign of the back-propagation gradients to produce targeted adversarial examples [19]:

$$x' = clip(x - \varepsilon \cdot sign(\frac{\partial L(x,t)}{\partial x})) \tag{6}$$

They showed that this approach can also fool non-linear models.

PGD An improved version of FGSM is the projected gradient descent (PGD) attack [33], which iteratively calculates a projection on the imperfect ε ball around the benign source image so that the adversarial example is a valid image and within the max distance limit:

$$x'_{i+1} = clip(x'_i - \alpha \cdot sign(\frac{\partial L(x'_i,t)}{\partial x'_i})) \tag{7}$$

where α controls how much perturbation would be made in each iteration and loosens the linear assumption of models for PGD. By default, PGD is configured to run for 40 iterations, and the step size of each iteration, α, is set to .01, where the values of each channel (of each pixel) in the images are normalized to $[0,1]$. The default loss function used in PGD is the CE loss.

Auto-PGD Croce and Hein were able to boost PGD's performance by intelligently setting its parameters (e.g., the step size α) [14]. Except for the number of iterations, the algorithm they proposed, Auto-PGD, does not require parameter tuning. Croce and Hein also showed that the loss function in PGD can markedly influence the attack's success rate. By default, Auto-PGD runs with the Difference of Logits Ratio (DLR) loss:

$$L_{DLR} = \frac{Z_t - Z_y}{Z_{\pi 1} - 0.5Z_{\pi 3} - 0.5Z_{\pi 4}} \tag{8}$$

where Z_πs are logits sorted from largest to smallest. Another commonly used loss function is the Carlini-Wagner (CW) loss [8]:

$$L_{CW} = -Z_t + max_{i \neq t} Z_i \tag{9}$$

CW loss and DLR loss both try to make the logit of the target class larger than other logits, which would cause the model to assign the highest probability to the target class.

2.4 Defenses

Researchers have proposed a variety of defenses to mitigate adversarial examples. Regarded as one of the strongest defenses [1], *adversarial training* augments the training process with correctly labeled adversarial examples to enhance models' robustness against them (e.g., [19, 29, 33, 46, 55]). As is common in related work [14, 17, 33, 54, 59], we evaluate our proposed attacks against adversarially trained neural networks (see Sec. 5.2).

Other defense strategies have also been explored. For example, different input transformations have been proposed to remove adversarial perturbations from inputs prior to classification (e.g., [21, 31, 43, 60]). Unfortunately, the majority of these defenses can be evaded by adaptive attacks that craft adversarial perturbations that survive the transformations [3, 4, 22]. Certain defenses attempt to detect adversarial examples (e.g., [18, 32, 34]). The majority of these defenses are unable to classify adversarial examples correctly, when detected. Moreover, researchers have also found that detection methods can often be evaded by adaptive attacks [4, 7]. Researchers have also investigated the use of formal methods to prove that the classification results of individual inputs are safe—i.e., no L_p-bounded perturbation can be added to the input to alter the classification results (e.g., [24, 48]). However, these methods have been shown to perform less well than adversarial training [5]. Finally, certified defenses provide provable accuracy guarantees on adversarial examples (e.g., [11, 30, 40, 56, 61]). However, these defenses are often effective for perturbations of smaller norms than adversarial training [30], or are ineffective against the threat models we study (e.g., high-dimensional perturbations with bounded L_∞-norms) [5, 26].

3 A Stronger Loss Function

In this section, we describe how we enhanced Auto-PGD's performance by improving its loss function. We first present an example scenario where a previous best-performing loss function falls short (Sec. 3.1). Then we describe a new loss function that mitigates the previous loss function's weaknesses (Sec. 3.2). We report on the performance of the new loss function in Sec. 6.

3.1 Observations on Previously Established Loss Functions

To produce a baseline for comparison, we ran Auto-PGD for 100 iterations, as it is used in previous work [14], with three loss functions: CW, DLR, and CE. We found that the CW loss performed the best against six out of seven CIFAR10 models as well as the ImageNet model we used. Fig. 2 shows our measurement results when using $\varepsilon = 16/255$ on CIFAR10, while results using different values of ε are shown in App. A.

Still, we identified specific instances in which Auto-PGD with CW loss failed to find adversarial examples. We identified that these failures can be often explained by the definition of CW loss. The CW loss tries to increase the difference between the logits of the target class and the highest among the logits of the non-target classes; however, there is no guarantee that the non-target class that has the highest logit, i (i.e., $Z_i = \max_{j \neq t}(Z_j)$), is the same class from iteration to iteration. Namely, decreasing Z_i may lead the logit of another non-target class to increase and become the highest in the next iteration. Thus, the target class might never have a chance to have the highest logit and the attack might never succeed.

Tab. 1 shows an example where Auto-PGD using CW loss fails to find an adversarial example. In that example, classes 3 and 4, neither of which is the target class, take turns to have highest confidence as the attack progresses. The target class, 0, never gets the highest confidence score as the loss function focuses on decreasing $\max_{j \neq t}(Z_j)$, which is the logit of one of class 3 or 4, while the logit of the other increases at the same time.

3.2 A New Loss Function

The example in Tab. 1 shows how a loss function that does not completely align with the attack's intent can sometimes fail to drive a candidate perturbation to become a successful adversarial input. Intuitively, the attacker's goal is to produce a perturbation for which the target class's logits and probabilities are higher than those of *all* other classes, even if by only a tiny amount, so that the target class could be selected as the prediction.

We capture this intuition through the design of a new loss function that contains a term that represents each logit. More precisely, the loss function is the sum over all such terms (each

Table 1: Confidence scores of three classes in five iterations from the process during which Auto-PGD with CW loss tries to perturb a CIFAR10 image. As explained in Sec. 2.1, relatively higher logits yield relatively higher confidence scores. In this example the adversary was trying to perturb an image from class 5 to class 0. The adversary tried to decrease the logit of class 3, Z_3, in iteration 81, causing the logit of class 4, Z_4, to increase. In the next iteration, the adversary decreased Z_4 but Z_3 increased. Z_3 and Z_4 took turns to have the highest logit and highest confidence score, and Auto-PGD with CW loss ultimately failed to find an adversarial example for this image.

Iteration Number	*Confidence*		
	Class 0 (target)	*Class 3*	*Class 4*
81	.1982	**.2072**	.2025
82	.1985	.1975	**.2146**
83	.1977	**.2077**	.2016
84	.1985	.1978	**.2140**
85	.1979	**.2082**	.2013

representing a logic), where each term's value is determined as follows.

- The term is zero if the logit is smaller than the logit of the target class. The adversary has no interest to further decrease this smaller logit, as it is already smaller than the highest logit and so does not influence the model's prediction.
- The term is positive if the logit does not belong to the target class and is higher than or equal to the logit of the target class. The adversary has to decrease all these non-target logits to make a successful perturbation. There are various ways to define this term, and we show one example below.

By minimizing this loss function, the adversary aims to decrease all the positive terms simultaneously and consequentially decrease all the logits that are higher than the logit of the target class in the same iteration.

Using the rectified linear unit (ReLU) function, we define this loss function, the Minimal Difference Loss (MD loss), as:

$$\sum_i ReLU(Z_i + \delta - Z_t) \tag{10}$$

where Z_t is the logit of the target class and Z_i are the logits of each class (as described in Sec. 2.1). δ is a minimal value introduced to mitigate the cases where non-target classes have equal logits with the logit of the target class. δ is set to $1e{-}15$ in the actual implementation (because $1.0-(1e{-}16)=1.0$) in Python due to finite arithmetic. Z_t is not excluded from Z_i as $ReLU(Z_t+\delta-Z_t)=\delta$ is always a constant and has no effect on the back-propagation gradients.

In contrast to CW loss, MD loss aims to decrease *all* logits that do not belong to the target class and are higher than the logit of the target class, rather than only the largest of the non-target-class logits. Logits that meet the following three requirements are more likely to decrease between iterations if the attacker is using MD loss instead of the CW loss:

- they do not belong to the target class;
- are higher than the logit of the target class; and
- are not the largest logit.

Hence, with MD loss the behavior demonstrated in Tab. 1 (non-target logits alternating at being the highest logit) is unlikely to occur and, intuitively, the attacker is more likely to succeed.

4 A Stronger Attack Method

In this section, we describe our new attack method, Constrained Gradient Descent (CGD). We start by explaining the intuition that drives the design and an enhancement (Secs. 4.1–4.3), and then detail our algorithm (Sec. 4.4). We report on the performance of the new attack in Sec. 6.

4.1 Learning to Stay Within the Distance Limit

PGD attacks, including Auto-PGD, enforce the L_∞ distance limit by executing a $clip(\cdot)$ computation in each iteration, eliminating any perturbations made outside the limit. However, the loss function used in PGD attacks does not take into account that $clip(\cdot)$ will be used. Hence, in PGD attacks, gradients, which are derivatives of loss functions against the current perturbation, may not accurately direct the perturbation. The gradients may push the attack toward perturbations outside the L_∞ distance limit, while successful perturbations within the L_∞ distance limit may lie in other directions.

Previous work has shown that attacks could *learn* to minimize their distance limit from the original benign image [8, 50]. We propose a new loss function that helps attacks learn to stay within a fixed L_∞ distance limit.

We include the L_∞ distance limit as part of the new loss function. This limit remains constant throughout the perturbation process. Before we start the perturbation, we know the upper and lower boundary of the value in channel $k \in \{0,1,2\}$ of the pixel with coordinates (i,j) in the current perturbation are

$$bound^{upper}_{i,j,k}(x) = \min(x_{i,j,k}+\varepsilon, 1) \tag{11}$$

and

$$bound^{lower}_{i,j,k}(x) = \max(x_{i,j,k}-\varepsilon, 0) \tag{12}$$

These two boundaries are fixed throughout the attack process. As $x_{i,j,k} \in [0,1]$, in each channel (of each pixel), the current

perturbation could at most go over one boundary at a time (for any one channel of any pixel). Thus, we can compute the distance, for each channel (of each pixel), by which the current perturbation x' goes over the boundary. We call this distance $Overrun_{i,j,k}(x')$ and define it as:

$$ReLU(x'_{i,j,k} - bound^{upper}_{i,j,k}(x)) + ReLU(bound^{lower}_{i,j,k}(x) - x'_{i,j,k}) \quad (13)$$

$Overrun_{i,j,k}(x')$ is 0 if x' stays within the ε-boundary for channel k of pixel (i,j).

We next add to the loss function the following term.

$$L_{boundary} = \sum_{i,j,k} (Overrun_{i,j,k}(x')^2) \quad (14)$$

$Overrun_{i,j,k}(x')$ is squared because when taking the derivative, $\frac{\partial L_{boundary}}{\partial x'_{i,j,k}}$ is proportional to $Overrun_{i,j,k}(x')$ and could direct the attack to learn to stay within the boundary.

Our new loss function is now:

$$w * L_{classification} + (1-w) * L_{boundary} \quad (15)$$

where $L_{classification}$ is MD-loss (as described in Sec. 3) and $w \in [0,1]$ is a weighting parameter. By decreasing w and weighting $L_{boundary}$ more heavily, we guide the attack to gradually move the perturbation in a direction that leads to the target class while not crossing the boundary, ultimately creating a valid adversarial example after rounding and clipping.

We also change the way that the attack iteratively updates the perturbation. To the best of our knowledge, previous L_∞ attacks that use gradients, $\frac{\partial Loss}{\partial x'_{i,j,k}}$, only use the sign of these gradients to generate adversarial examples (e.g., [14, 17, 19, 27, 28, 33, 35, 53]), while we believe the magnitude of the gradients also conveys information about the amount by which it is helpful to change a channel (of a pixel).

Leveraging the magnitude of the gradients also helps avoid artificially (and unhelpfully) large step sizes that would make the perturbations step over the boundary, which we observed can happen in, e.g., PGD attacks. In order to reduce fluctuation, we feed the gradients to Adam optimizers [25]: Each channel (of each pixel) of the image has an Adam optimizer, and we add the outputs of the optimizers to the current perturbation to generate a new perturbation.

4.2 Driving out of Local Minima

While $L_{boundary}$ accounts for any channels (of any pixels) potentially crossing the ε-boundary, we observed that the attack could become trapped in a local minimum of $L_{boundary}$ that occurs when most channels (of most pixels) are within the boundary but a small number is far beyond the boundary. To prevent this, we set a threshold distance outside the upper and lower ε boundaries. The threshold gradually decays as the attack progresses. Model-specific constants, namely pre-defined fixed ratios and checkpoints, control where the threshold starts and how fast it decays. If the current perturbation is outside the threshold, we halve the weight w of $L_{classification}$ to urge the attack to focus more on staying within the boundary.

We ran grid searches to choose the starting threshold and decay interval for each defense. We examined starting thresholds $\in [.25\varepsilon, 10\varepsilon]$, starting the weighting parameter of the loss function $w \in [0.01, 0.5]$, and decay intervals $\in [5, 30]$. For all defenses, the optimal starting weighting parameter w was 0.1 and the optimal decay interval was every 15 iterations. The optimal starting threshold was 8ε for SIE+20, 4.75ε for DSL+20, and 1.75ε for WRK20, HLM19, WZY+20, SWM+20, CRS+19, and WXW20.

4.3 Smart Rounding

As we described in Sec. 2.2, we always make a 8-bit RGB format copy of the current perturbed image for the models to predict on. However, we found arithmetic rounding might not be the optimal approach for the adversary to generate such a copy. As white-box attacks, including the PGD attacks, compute $\frac{\partial L(x'_i,t)}{\partial x'_i}$ in each iteration, the adversary can use this information to perform *smart rounding* to potentially find more adversarial examples. We implement two instances of smart rounding per iteration. In particular, we make two 8-bit RGB format copies of the generated continuous domain current perturbation x'_{i+1} and round them up in the direction pointed to by the gradients as follows.

$$x_{SR1} = round(x'_{i+1} * 255 + sign(\frac{\partial L(x'_i,t)}{\partial x'_i}) * .499999)/255 \quad (16)$$

$$x_{SR2} = round(x'_{i+1} * 255 + sign(\frac{\partial L(x'_{i+1},t)}{\partial x'_{i+1}}) * .499999)/255 \quad (17)$$

We take the predictions of the model on the clipped perturbations $clip(x_{SR1})$ and $clip(x_{SR2})$ and if either of them matches the target class, we stop perturbing this example and declare the attack a success. If neither of them matches, the attack will continue with x'_{i+1}, which is still in the continuous domain. An attack with smart rounding yields a success rate that is no worse than that of its counterpart without smart rounding, as smart rounding consists only of two additional evaluations in each iteration and the attack process in the continuous domain remains unchanged.

CGD requires smart rounding to outperform Auto-PGD with MD loss, either with or without smart rounding. This is because CGD attacks are designed to explore the ε boundary from the outside, and so the candidate perturbation often needs to be nudged to be within the ε boundary.

4.4 The Constrained Gradient Descent Algorithm

Combining all the above, we define a new attack: *Constrained Gradient Descent*. Fig. 1 illustrates an example path of CGD where a candidate perturbation learns to stay within the L_∞ distance limit.

We start the attack with a random initial perturbation to better explore the space of possible adversarial examples, as this was shown to be helpful in prior work [13, 33, 37]. In stage 1, we move each channel (of each pixel) by ε in the direction of the gradients. This is a quick way to take a substantial step in the direction of the target class. In stage 2, the candidate perturbation continues to move toward the target class, and potentially moves outside the threshold, as the loss function is dominated by $L_{classification}$. If the candidate perturbation moves beyond the threshold, the algorithm moves to stage 3: since the candidate perturbation is outside the threshold, $L_{boundary}$ is more heavily weighted, which pushes the candidate perturbation back inside the threshold. After a fixed number of iterations the algorithm enters stage 4, in which the threshold itself moves toward the ε boundary, which also forces the candidate perturbation to move closer to the ε boundary. The number of iterations in stages 2–4 is not fixed, and an attack could succeed in any stage. We observe that most perturbations fluctuate in their L_∞ distance from the source image as the attack tries to balance between $L_{classification}$ and $L_{boundary}$.

Algorithm 1 presents the attack's pseudocode. The inputs to the algorithm are: the benign example x; the $model(\cdot)$ function that emits the logits for given inputs; the number of iterations $n_{iterations}$; the ε distance; a set of constants tuned per model, *threshold* and *checkpoints*; and the target class *target*. We allow the tuning of *threshold* and *checkpoints* per model because we are running white-box attacks and hence the adversary has the freedom to choose the attack's parameters per model. All other parameters do not require tuning per model. In line 1, we add a random initial perturbation to the benign sample, moving it to the ε boundary. In lines 3 and 4 of the algorithm, before we start the iterations, we compute $bound^{upper}(x)$ and $bound^{lower}(x)$. From line 6 to 8, we adjust the threshold based on pre-set constants. We compute the logits Z of the model regarding the current perturbation x' in line 9 and compute the MD loss in line 10. Then we compute the *Overrun* in line 11 and the $L_{boundary}$ in line 12, which motivates the adversarial example to move to be within the L_∞ distance limit. From line 13 to line 15, we adjust the weight w, and then use this weight to compute the total loss as a weighted sum in line 16. We compute the gradients from the loss function in line 17 and apply changes from line 22 to line 27. In the first iteration, the candidate perturbation is pushed by a distance of ε to more quickly reach the boundary. In later iterations, the Adam optimizer is used to reduce fluctuation. From line 28 to 31 (highlighted in pink), we perform a nor-

Input : x, $model(\cdot)$, $n_{iterations}$, ε, *threshold*, *checkpoints*, *target*
Output : x'

Function *Sround* $(x', gradients)$ **:**
  return $clip(round(x' * 255 + sign(\frac{\partial L(x',t)}{\partial x'}) * .499999)/255)$;

Function *Nround* (x') **:**
  return $clip(round(x' * 255)/255)$;

1: $x' \leftarrow clip(x + 2 * rand(x.shape) - 1)$;
2: $w \leftarrow 0.1$;
3: $bound^{upper}_{i,j,k}(x) \leftarrow min(1, x_{i,j,k} + \varepsilon)$;
4: $bound^{lower}_{i,j,k}(x) \leftarrow max(0, x_{i,j,k} - \varepsilon)$;
5: **for** $iteration \leftarrow 1$ **to** $n_{iterations}$ **do**
6: **if** $iteration \in checkpoints$ **then**
7: $threshold \leftarrow threshold/2$;
8: **end if**
9: $Z \leftarrow model(x')$;
10: $L_{classification} \leftarrow \sum_i ReLU(Z_i + \delta - Z_t)$;
11: $Overrun_{i,j,k}(x') \leftarrow ReLU(x'_{i,j,k} - bound^{upper}_{i,j,k}(x)) + ReLU(bound^{lower}_{i,j,k}(x) - x'_{i,j,k})$;
12: $L_{boundary} \leftarrow \sum_{i,j,k}(Overrun_{i,j,k}(x')^2)$;
13: **if** $\exists i,j,k \in Overrun_{i,j,k}(x') > threshold$ **then**
14: $w \leftarrow w/2$;
15: **end if**
16: $loss \leftarrow w * L_{classification} + (1-w) * L_{boundary}$;
17: $gradients \leftarrow \frac{\partial loss}{\partial x'}$;
18: $x_{test} \leftarrow Sround(x', gradients)$;
19: **if** $argmax(model(x_{test})) == target$ **then**
20: **return** x_{test};
21: **end if**
22: **if** $iteration == 1$ **then**
23: $x' \leftarrow clip(x' - \varepsilon * sign(gradients))$
24: **else**
25: $changes \leftarrow Adamoptimizer(gradients)$;
26: $x' \leftarrow x' - changes$;
27: **end if**
28: $x_{test} \leftarrow Nround(x')$;
29: **if** $argmax(model(x_{test})) == target$ **then**
30: **return** x_{test};
31: **end if**
32: $x_{test} \leftarrow Sround(x', gradients)$;
33: **if** $argmax(model(x_{test})) == target$ **then**
34: **return** x_{test};
35: **end if**
36: **end for**
37: **return** *failed attack*;

Algorithm 1: CGD

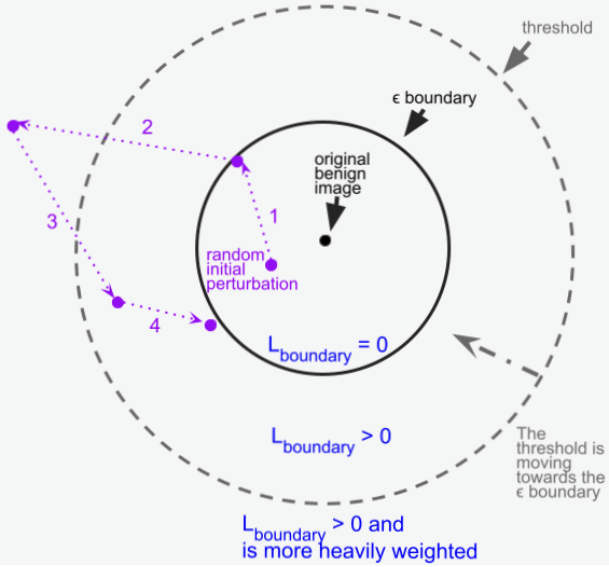


Figure 1: This is an example path of a CGD attack with a L_∞ distance limit. We start with a random initial perturbation. In stage 1, we push the current perturbation to the ε boundary. In stage 2, the current perturbation moves beyond the threshold and $L_{boundary} > 0$. In stage 3, the current perturbation is pushed inside the threshold as $L_{boundary}$ is more heavily weighted. In stage 4, the current perturbation moves closer to the ε boundary, as does the threshold.

mal rounding check, calling the normal rounding subroutine $Nround(\cdot)$. From line 18 to 21 and line 32 to 35 (highlighted in yellow) we perform two smart rounding checks, calling the smart rounding subroutine $Sround(\cdot)$. These rounding checks enforce the successful adversarial example to be inside the L_∞ distance limit.

Both rounding routines (which check whether the current perturbation is a successful attack) clip the image so that it is within the ε boundary. When the candidate perturbation is very close to the ε boundary, clipping has minimal effects on the logits, $L_{classification}$, and probabilities assigned to classes, meaning that an image just outside the ε boundry is rounded to become both within the ε boundary and potentially a successful attack.

5 Evaluation Setup

In this section, we introduce how we set up experiments to compare attacks. Factors other than the algorithm, such as the random initialization chosen [51] and which classes are targeted can also influence attacks' performance. Thus, before we introduce our results (Sec. 6), we explain how we executed reproducible attacks to enable meaningful and fair comparisons.

5.1 Datasets

In this work, we used the CIFAR10 and ImageNet datasets—two standard datasets that are commonly used for classification tasks. Both datasets contain colored images of objects in 8-bit RGB format as described in Sec. 2.2. Each image in CIFAR10 has 32×32 pixels, whereas we resized each image in ImageNet to have 224×224 pixels. We evaluated attacks on the test set of the datasets. CIFAR10 has 10,000 images in its test set, and ImageNet has 50,000 images in its test set.

5.2 Benchmarks

As explained in Sec. 2.4, we evaluated attacks using models trained via adversarial training, as this is a common benchmark in related work [14, 17, 33, 54, 59], and is regarded as a strong defense [1]. Specifically, we used seven pre-established adversarially trained models for CIFAR10: CRS+19 [9], DSL+20 [16], HLM19 [23], SWM+20 [44], WRK20 [57], WXW20 [58] and WZY+20 [55]. All of these models were trained with $\varepsilon = 8/255$ and are publicly available via `RobustBench` [12], a standard library for robustness evaluation of neural networks. We selected the models using the following criteria: they rank highly on robustness in the public `RobustBench` evaluation; they have a PyTorch implementation; and they fit into the 11GB memory of our NVIDIA GeForce RTX 2080 GPUs. These models have have also been found to be robust in evaluations other than `RobustBench`'s [20, 41]. The seven models can be categorized into four non-exclusive groups: some propose surrogate loss functions [16, 55, 58], some use pre-training or semi-supervised learning techniques [9, 23, 55], some try to wisely tune the training process [57], and some apply pruning to their models [44]. Similarly to prior work (e.g., [49, 57]), we evaluated attacks against these models with $\varepsilon = 8/255$ and $16/255$. We also found two versions of SIE+20 [42], pre-established and publicly available adversarially trained models on ImageNet—one adversarially trained with $\varepsilon = 4/255$ and another with $\varepsilon = 8/255$. We evaluated attacks against each version of this model with the same ε with which it was trained.

5.3 Experiment Setup

Croce et al. found that PGD attacks find more adversarial examples the more iterations they run [14]. Auto-PGD declares the number of iterations as its only parameter. In this work, we ran all attacks for 100 iterations—the default configuration of Auto-PGD [14]—to fairly compare the attack methods.

We observed that some PyTorch functions could yield unreproducible results in different runs. To keep our measurement results reproducible, we set `torch.backends.cudnn.benchmark` to `False` and `torch.backends.cudnn.deterministic` to `True`. In addition, we found that PGD attacks, including Auto-PGD,

always start with a random perturbation and this could slightly influence the result. Hence, we ran these attacks multiple times, each time with a different random seed. We also fixed the batch size so that with the same random seed we got the same random initial perturbation. We used a batch size of 512 images for CIFAR10 and a batch size of 10 images for ImageNet. In addition to the random initialization, we also picked a random target class for each image in the testing sets of datasets which we used to evaluate attacks. Due to limited computing resources, we were not able to run attacks targeting every incorrect class, especially for ImageNet which has 1,000 classes. The target class was intentionally selected to be different from the label, the correct class of the image. We chose the difference *target_offset* between the target class and the correct class, using the following formula:

$$target_offset_i = floor(rand() * (N_{classes} - 1)) + 1 \quad (18)$$

where i is the index of images in the testing set, $i \in [0, 10000)$ for CIFAR10 and $i \in [0, 50000)$ for ImageNet, and $N_{classes}$ is the number of classes, 10 for CIFAR10 and 1,000 for ImageNet. The $rand()$ function generates a float $\in [0, 1)$. The target class was

$$t_i = (y_i + target_offset_i) \bmod N_{classes} \quad (19)$$

for $i \in [0, N_{images})$. We only used one random seed (specifically, 0) for the *target_offset* in CIFAR10, as on average there are $10000/(9*10)$ images in each source-target class pair, whereas we used five random seeds (specifically, 0–4) for the *target_offset* of the ImageNet dataset, as 50,000 images cannot cover all the 999×1000 source-target class pairs. We also observed that some class pairs in ImageNet (e.g., "African crocodile" and "American alligator") are closer than other class pairs (e.g., "African crocodile" and "thunder snake"), and are significantly more easy to perturb into each other. Random numbers were generated as a vector of length 10,000 for CIFAR10 and 50,000 for ImageNet. Hence, the *target_offset* was the same for the same image with the same random seed, but *target_offsets* was not the same across all images when we used the same random seed.

Overall, for CIFAR10, we measured the success rate against seven defenses, using 20 random initializations and two ε values per image, resulting in a total of 280 sets of 10,000 attack attempts. In the case of ImageNet, we used five random initializations, five target offsets, and two ε values per image, thus resulting in 50 sets of 50,000 attack attempts.

6 Evaluation Results

In this section, following the setup described in Sec. 5, we compare Auto-PGD using our MD loss with Auto-PGD using previously established loss functions and also our CGD attack with Auto-PGD. We first report on raw results (Sec. 6.1) and then on the statistical tests we performed (Sec. 6.2) to demonstrate that CGD outperformed the previously best Auto-PGD with statistical significance. We also examine the degree to which different attack methods succeed and fail on the same source images (Sec. 6.3).

6.1 Raw Results

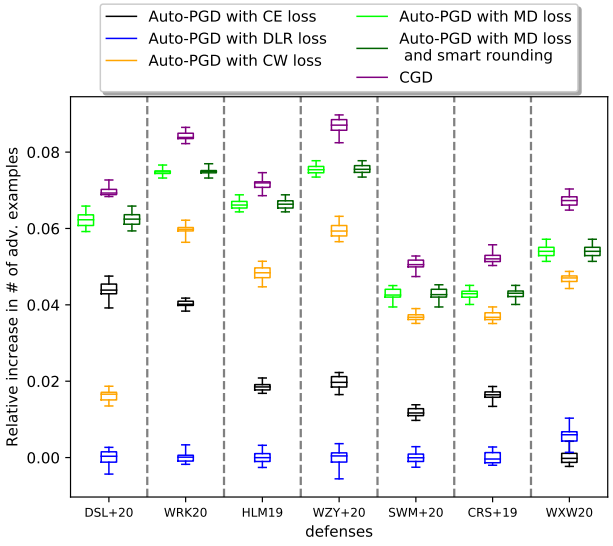


Figure 2: This figure shows the relative improvement in the number of adversarial examples found by attacks on different defenses compared to the worst-performing attack. Experiments were performed using 10,000 images from the test set of CIFAR10, $\varepsilon = 16/255$, 20 different random initial perturbations, and a fixed random target offset. The worst-performing attack for each defense (with a median of 0) was selected as the baseline. The y-axis denotes the improvement compared to the average performance of the baseline. For example, 0.08 on the y-axis indicates 8% more adversarial examples found compared to the baseline.

As described in Sec. 5.3, we made 280 sets of 10,000 attack attempts on CIFAR10 and 50 sets of 50,000 attack attempts on ImageNet. We compared our two improvements, Auto-PGD using our MD loss (Sec. 3) and CGD (Sec. 4), to Auto-PGD using three pre-established loss functions: CE loss, DLR loss, CW loss. Implementations of Auto-PGD with CE loss and DLR loss are the ones published by the authors [14]. We perform this comparison on two datasets and multiple defenses and values of ε as described in Sec. 5. For a more complete comparison with CGD, we also ran Auto-PGD using our MD loss and smart rounding(Sec. 4.3).

On average, Auto-PGD with MD loss found more adversarial examples than Auto-PGD with any of the three other loss functions. Auto-PGD with MD loss and smart rounding performed slightly better than Auto-PGD with MD loss but without smart rounding. CGD performed better than Auto-

PGD with both MD loss and smart rounding. The ranking of attacks other than Auto-PGD with MD loss and CGD varied slightly depending on the defense.

Fig. 2 shows the relative number of adversarial examples these attack methods found on CIFAR10 with $\varepsilon = 16/255$. Out of 10,000 attempts, the baseline, worst-performing attack succeeded on average 5,998 times on DSL+20, 5,881 times on WRK20, 4,482 times on HLM19, 4,669 times on WZY+20, 4,646 times on SWM+20, 4,607 times on CRS+19, and 3,799 times on WXW20. Which attack performed least well varied by defense. The best of the previously proposed attacks—Auto-PGD with one of CE, DLR, or CW loss—was on average 0.4–5.4% better than the baseline attacks. Which loss function led to best performance depended on the defense. Auto-PGD with MD loss was on average 0.6–4.5% better than the best performing Auto-PGD that did not use MD loss; and CGD was on average 1.3–5.3% better than the best performing Auto-PGD that did not use MD loss. CGD outperformed Auto-PGD with both MD loss and smart rounding in 139 out of 140 sets of attempts and outperformed Auto-PGD (with any other loss function) in all of the sets of attempts. Depending on the defense, the ranking of the other attack methods changed, but CGD and Auto-PGD with MD loss always performed better than other attack methods, as shown in Fig. 2.

We observed similar results with different values of ε and when using the ImageNet dataset, as shown in App. A. When evaluated against SIE+20 on the ImageNet dataset with $\varepsilon = 4/255$, CGD was 10.9% better than Auto-PGD with CW loss (the previous best Auto-PGD). When using $\varepsilon = 8/255$, CGD was 13.7% better than Auto-PGD with CW loss against SIE+20 on ImageNet.

Tab. 2 shows the proportions of sets of attempts where smart rounding increases the success rate of Auto-PGD with MD loss. In some sets of attempts, Auto-PGD with MD loss performed exactly the same regardless of smart rounding; whereas in sets of attempts where smart rounding did improve the performance, the improvement is always less than 0.1% on CIFAR10 and less than 0.4% on ImageNet. Smart rounding does not significantly improve the performance of Auto-PGD with MD loss.

Table 2: This table shows the percentage of sets of attempts where smart rounding improved the performance of Auto-PGD with MD loss. CIFAR10 has 7 defenses $\times$ 20 random initial perturbations = 140 sets of attempts for each ε, whereas ImageNet has 5 random initial perturbations $\times$ 5 target offsets = 25 sets of attempts for each ε.

ε	*Dataset*	
	CIFAR10	*ImageNet*
4/255	-	60%
8/255	26%	88%
16/255	40%	-

6.2 Statistical Analysis

We also performed statistical analysis to compare the performances of different attack methods. We defined a variable $CondtionalSuccess_i$ for an image i as the total number of successful perturbations made by attack A with all the defenses, random initial perturbations, and all the random *target_offset*s we tried given the specific dataset and ε. $CondtionalSuccess_i \in [0, 140]$ on CIFAR10 while $CondtionalSuccess_i \in [0, 25]$ on ImageNet. Each $CondtionalSuccess_i$ is independent. We used the Wilcoxon signed rank test to compare $CondtionalSuccess_i$ of CGD and Auto-PGD with MD loss and smart rounding. The Wilcoxon signed rank test assumes the two attack methods perform differently on every image. In our case, the attacks perform equally well on some images. Previous work suggests a mitigation to enable using the Wilcoxon signed rank test when the two attack methods perform equally well on some images: using an adjusted normal approximation of the test statistic of the Wilcoxon signed rank test instead of using the standard test statistic to compute the p-value [15]. This approach requires more than 25 independent trials; our data satisfies this requirement.

We performed the one-sided Wilcoxon signed rank test with the null hypotheses that Auto-PGD with MD loss and smart rounding had equal or better performance than CGD for each combination of ε, dataset, and defense that we tried. Overall, we conducted 16 statistical tests, for the 16 different combinations we had. Thus, to account for the multiple tests, we used Bonferroni correction to adjust the confidence level α to $.05/16 = 0.003125$.

The result is shown as Tab. 3. The p-values are below α in 13 out of 16 tests. Namely, CGD performed statistically significantly better than Auto-PGD with MD loss across 13 combinations of ε, dataset, and defense that we tried.

We performed a similar one-sided Wilcoxon signed rank test with the null hypotheses that the best performing attack among Auto-PGD using the DLR loss, CW loss, and CE loss performed equal to or better than CGD in each combination of value of ε, dataset, and defense that we tried. Again, we used the adjusted normal approximation of the test statistic and Bonferroni corrections. All the p-values were far smaller than $\alpha = .05/16 = 0.003125$, and so we reject the null hypotheses in all 16 cases, hence demonstrating that CGD significantly outperformed Auto-PGD with the losses proposed in prior work across each combination of ε, dataset, and defense that we tried.

6.3 Uniqueness of Attacks

As we observed that CGD found more adversarial examples than other attack methods did in most of the sets of attempts, we were wondering if the adversarial examples other attacks found could be a subset of those CGD found. However, we discovered that among the attack methods we tried, each of

Table 3: This table shows the result of the Wilcoxon signed rank test with an adjusted normal approximation for the null hypotheses that Auto-PGD using the MD loss and smart rounding performed equal or better than CGD in each combination of value of ε, dataset, and defense that we tried. We follow the common practice to include the Wilcoxon statistics along with the p-values. Test results where p-values are smaller than α are shown in bold. We reject the null hypothesis in 13 of the 16 tests.

dataset	*defense*	*L_∞ distance limit*			
		$\varepsilon = 8/255$		*$\varepsilon = 16/255$*	
		W statistic	*p value*	*W statistic*	*p value*
CIFAR10	DSL+20 [16]	**2562673.5**	**8e-13**	5226620.0	0.134
	WRK20 [57]	**508542.0**	**4e-7**	**2416685.0**	**3e-7**
	HLM19 [23]	229636.5	0.031	**2533997.5**	**6e-5**
	WZY+20 [55]	289324.0	0.026	**3092694.0**	**1e-7**
	SWM+20 [44]	**269493.5**	**0.0026**	**2298465.0**	**3e-7**
	CRS+19 [9]	**299411.0**	**8e-14**	**2606579.0**	**3e-9**
	WXW20 [58]	**249655.0**	**1e-4**	**1818443.5**	**6e-13**
ImageNet		*$\varepsilon = 4/255$*		*$\varepsilon = 8/255$*	
		W statistic	*p value*	*W statistic*	*p value*
	SIE+20 [42]	**9172979.0**	**3e-9**	**21150537.0**	**6e-13**

Table 4: This table shows the result of the Wilcoxon signed rank tests with an adjusted normal approximation for the null hypotheses that the best performance among Auto-PGD using the DLR loss, CW loss and CE loss performed equal or better than the performance of CGD on each image across all values of ε, datasets, and defenses. We follow the common practice to include the Wilcoxon statistics along with the p-values. Test results where p-values are smaller than α are shown in bold. We reject the null hypothesis in all 16 tests.

dataset	*defense*	*L_∞ distance limit*			
		$\varepsilon = 8/255$		*$\varepsilon = 16/255$*	
		W statistic	*p value*	*W statistic*	*p value*
CIFAR10	DSL+20 [16]	**2408287.0**	**1e-8**	**6636015.5**	**7e-10**
	WRK20 [57]	**657802.0**	**7e-9**	**3588117.5**	**7e-27**
	HLM19 [23]	**528566.0**	**2e-7**	**3331182.5**	**4e-16**
	WZY+20 [55]	**508543.5**	**5e-6**	**3923168.5**	**2e-18**
	SWM+20 [44]	**508660.5**	**1e-8**	**2523202.5**	**4e-10**
	CRS+19 [9]	**508622.0**	**2e-7**	**2849349.5**	**4e-11**
	WXW20 [58]	**498728.0**	**6e-9**	**2014654.0**	**2e-16**
ImageNet		*$\varepsilon = 4/255$*		*$\varepsilon = 8/255$*	
		W statistic	*p value*	*W statistic*	*p value*
	SIE+20 [42]	**19071169.5**	**4e-57**	**55335060.5**	**6e-171**

them found a slightly different set of successful adversarial examples. Fig. 3 shows that when using $\varepsilon = 16/255$ against DSL+20 on CIFAR10, each attack found some adversarial examples which another method did not; Fig. 4 shows that when using $\varepsilon = 16/255$ against each of the defenses we used on CIFAR10, each attack found some adversarial examples which any of the other methods did not. We observed the same phenomenon when we used a different value of ε or evaluate against a different defense (detailed results are shown in App. B).

At the same time, when the same attack method is given different random initial perturbations, the attack always found a slightly different set of adversarial examples, within the 20 random initial perturbations we tried. For example, Fig. 5 shows that when using $\varepsilon = 16/255$ against DSL+20 on CIFAR10, running the same attack with more random initializations kept finding new adversarial examples. This also held for other values of ε and against other defenses, as shown in App. C.

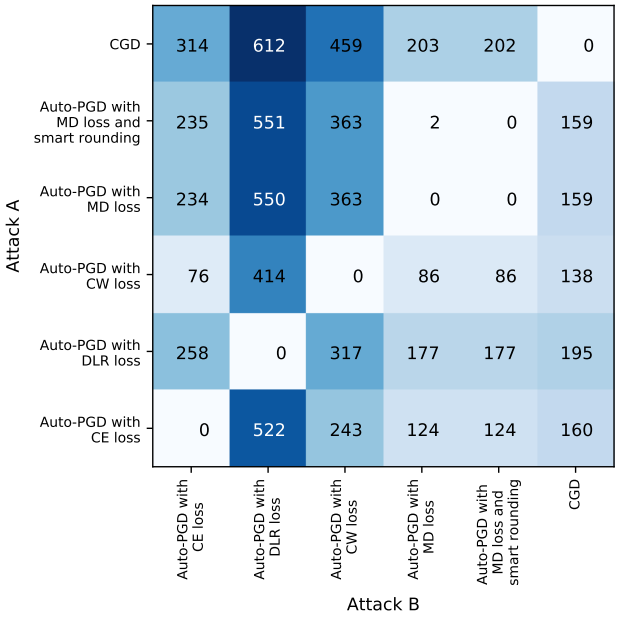


Figure 3: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 16/255$ against the DSL+20 [16] defense, no decimals are kept. Compared to other methods, each attack method found some unique adversarial examples. The only exception is between Auto-PGD using MD loss, with and without smart rounding, as the difference is additional tests in each iteration. The adversarial examples that Auto-PGD with MD loss but without smart rounding found comprise a subset, although not necessarily a proper subset, of the adversarial examples that Auto-PGD with MD loss and smart rounding found.

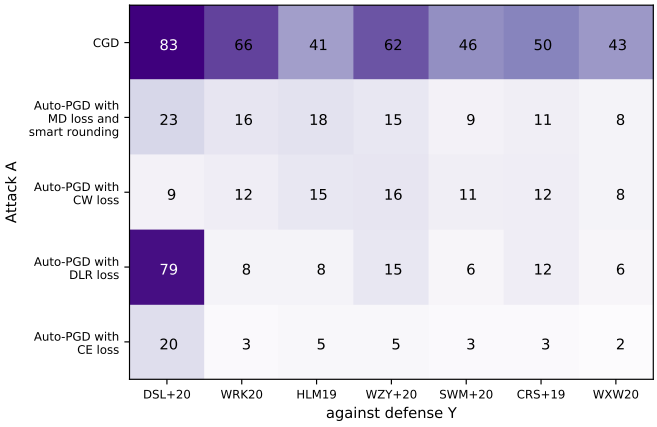


Figure 4: This is the average number of successful adversarial examples found by attack A but not by any other attack methods, against each defense Y in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 16/255$, no decimals are kept. Each of these attack methods found some unique adversarial examples. Auto-PGD using MD loss but without smart rounding is not included in this plot as the adversarial examples it found are a subset, although not necessarily a proper subset, of those found by Auto-PGD using MD loss and with smart rounding.

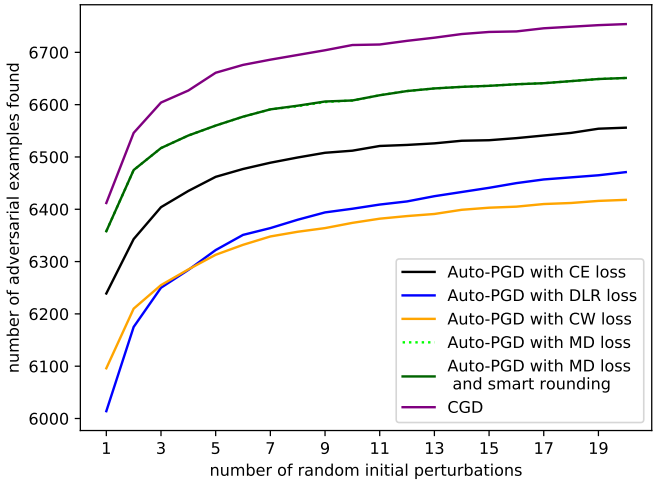


Figure 5: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 16/255$, against the DSL+20 defense. Each attack method finds more adversarial examples when they are allowed to use more different initial perturbations, up to 20 that we have tried

7 Discussion

In this section, we discuss the potential to use CGD as a *framework* for attacks (Sec. 7.1), e.g., with different distance metrics and loss functions. We describe and evaluate one such use, in which we instantiate CGD for untargetted attacks (Sec. 7.2).

7.1 CGD as a Framework

In Sec. 6 we showed that CGD outperformed the previous best attack in targeted tasks, with substantial difference and statistical significance. Similarly to how Auto-PGD is a member of the PGD family attacks, the specific attack we explore in this paper could be viewed as member of a broader CGD family. Auto-PGD improves the performance of PGD attacks by using alternative loss functions and wisely tuning parameters; similar tweaking could also apply to CGD attacks: using alternative loss functions and wisely tuning parameters could yield a stronger attack within the CGD family, as they better match the goal of the adversary. Meanwhile, as the loss function of CGD can be separated into two components, $L_{boundary}$ and $L_{classification}$, each of those could be tuned. We demonstrated that CGD implemented with specific loss functions and parameters outperformed previous best attack; other variants of CGD could perform better yet. Our work opens the door for future work to find stronger attacks using CGD as a framework. CGD might also be extended to other distance metrics besides L_∞ and also to untargeted attacks.

7.2 Applying CGD to Untargeted Attacks

We demonstrated that CGD outperformed the previous best white-box targeted attacks in Sec. 6. Here we explore an untargeted variant of CGD as an example of extending CGD to other types of attacks, and, specifically, how other loss functions can be incorporated in the general CGD approach.

Previous works find untargeted adversarial examples by using one of several loss functions. One example is cross entropy loss [14, 19, 33], which we described in 2.1: $L_{CE} = -\log P_y = -Z_y + \log(\sum_{j=1}^{K} e^{Z_j})$. Other work proposed the untargeted attack version of the Difference of Logits Ratio loss (DLR loss) [14]:

$$L_{DLR} = \frac{Z_y - max_{i \neq y} Z_i}{Z_{\pi 1} - Z_{\pi 3}} \tag{20}$$

and showed that Auto-PGD with L_{DLR} finds more adversarial examples than Auto-PGD with L_{CE} [14]. Yet other work uses the untargeted version of CW loss [8]:

$$L_{CW} = -Z_y + max_{i \neq y} Z_i \tag{21}$$

Previous works show that by iteratively maximizing these loss functions the adversary can find untargeted adversarial examples [8, 14].

Because the concept of staying within the L_∞ distance limit is the same in targeted and untargeted attacks, the loss component for capturing the task of staying within the L_∞ distance limit, $L_{classification}$, is the same as described in Sec. 4.

$L_{classification}$ decreases to 0 if and only if the adversarial example is within the L_∞ distance limit. Hence, we design the loss component that captures the task of forcing misclassification, such as the MD loss, to also decrease to 0 when the adversary succeeds. The purpose of this design is to have the overall loss function, which is a weighted sum of the two components, decrease to 0 if and only both two tasks have been both successfully completed. To create the loss component that captures the task of avoiding correct classification, we define a variant of the CW loss as follows:

$$L_{CW^*} = ReLU(Z_y + \delta - max_{i \neq y} Z_i) \tag{22}$$

where δ is a minimal value, set to $1e{-}15$ (as with Sec. 3.2). We replace MD loss in Alg. 1 with L_{CW^*} as $L_{classification}$ and hence obtain the untargeted variant of the CGD attack, $\text{CGD}_{\text{untarg}}$.

Minimizing L_{CW^*} and maximizing L_{CW} achieve exactly the same result: when $Z_y \geq max_{i \neq y} Z_i$, both loss functions minimize $Z_y - max_{i \neq y} Z_i$; when $Z_y < max_{i \neq y} Z_i$, the untargeted attack has already succeeded. Hence, an Auto-PGD that minimizes L_{CW^*} and an Auto-PGD that maximizes L_{CW} behave exactly the same on each image, and we do not report the performance of Auto-PGD that uses L_{CW} alone.

We evaluated this attack on the CIFAR10 dataset. We use the L_∞ distance metric for untargeted attacks as we did for targeted attacks. We use three values of ε: $4/255$, $8/255$, and $16/255$. We use the same seven defenses as previously [9, 16, 23, 44, 55, 57, 58], again trained with $\varepsilon = 8/255$ from robust-bench [12], as benchmarks to evaluate untargeted attacks. We evaluated untargeted attacks by the number of adversarial examples they find within the L_∞ distance limit. We ran Auto-PGD, again with 100 iterations as in its default configuration, with L_{DLR} and L_{CW^*}. We also ran $\text{CGD}_{\text{untarg}}$ and Auto-PGD with L_{CW^*} and smart rounding for the same number of iterations. We ran each attack with five different random initial perturbations, using seeds 0–4, and a batch size of 512.

On average $\text{CGD}_{\text{untarg}}$ outperformed Auto-PGD with L_{DLR} and Auto-PGD with L_{CW^*}, with or without smart rounding. Fig. 6 shows the result for $\varepsilon = 8./255$, where $\text{CGD}_{\text{untarg}}$ outperformed the next best method, Auto-PGD with L_{CW^*} and smart rounding in 34 out of the 35 sets of attempts. We observe similar results with other values of ε (shown in App. D).

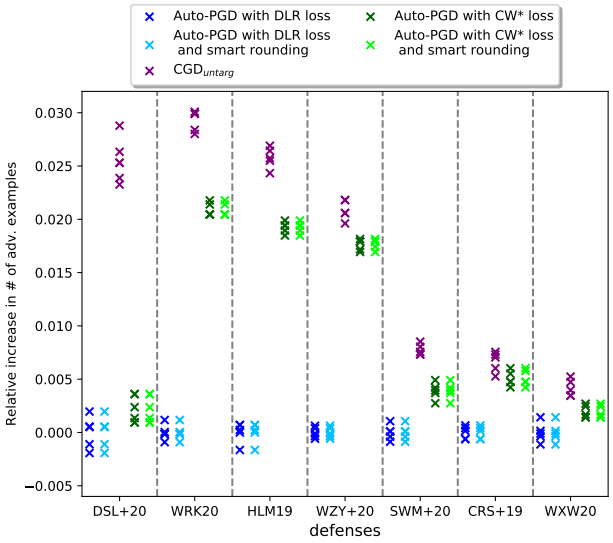


Figure 6: This figure shows the relative improvement in the number of adversarial examples found by untargeted attacks on different defenses compared to the worst-performing attack. Experiments were performed using 10,000 images from the testing set of CIFAR10. We ran attacks using $\varepsilon = 8/255$, and five different random initial perturbations with seeds 0–4. The result is normalized by the mean of the worst performing method against each model. On average, $\text{CGD}_{\text{untarg}}$ performed better than Auto-PGD with L_{DLR} and Auto-PGD with L_{CW^*}, with or without smart rounding.

8 Conclusion

In this work we improved a previously established white-box, targeted evasion attack by using a new loss function. We also proposed a yet stronger attack that learns to approach and explore the ε-boundary. We demonstrated the efficacy of both the new loss function and the new attack on two datasets (CIFAR10 and ImageNet), for multiple values of ε, and against multiple defenses; in all cases, our methods outperformed the best of the attacks we compared against. Although our attacks outperformed others, we also found that each attack we tested can succeed on *some* inputs on which other attacks fail. Finally, we showed how use our new attack method as a general framework for attacks and demonstrated its utility by instantiating it into a stronger *untargeted* attack.

References

[1] Naveed Akhtar, Ajmal S. Mian, Navid Kardan, and Mubarak Shah. Advances in adversarial attacks and defenses in computer vision: A survey. *ArXiv*, 2021. https://arxiv.org/abs/2108.00401.

[2] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: a query-efficient black-box adversarial attack via random search. In *ECCV*, 2020.

[3] Anish Athalye and Nicholas Carlini. On the robustness of the cvpr 2018 white-box adversarial example defenses. *ArXiv*, abs/1804.03286, 2018.

[4] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security:

Circumventing defenses to adversarial examples. In *International conference on machine learning*, 2018. https://arxiv.org/abs/1802.00420.

[5] Suman Jana Bai Li, Shiqi Wang and Lawrence Carin. Towards understanding fast adversarial training. *ArXiv*, abs/2006.03089, 2020.

[6] Battista Biggio, Igino Corona, Davide Maiorca, Blaine Nelson, Nedim Srndic, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 2013. https://arxiv.org/abs/1708.06131.

[7] Nicholas Carlini and David Wagner. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM workshop on artificial intelligence and security (AISec)*, 2017.

[8] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy*, volume 1, pages 39–57, 2017. https://arxiv.org/abs/1608.04644.

[9] Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, Percy Liang, and John C. Duchi. Unlabeled data improves adversarial robustness. In *Conference on Neural Information Processing Systems*, 2019. https://arxiv.org/abs/1905.13736.

[10] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, 2017.

[11] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, 2019.

[12] Francesco Croce, Maksym Andriushchenko, Vikash Sehwag, Edoardo Debenedetti, Nicolas Flammarion, Mung Chiang, Prateek Mittal, and Matthias Hein. Robustbench: a standardized adversarial robustness benchmark. arXiv preprint 2010.09670, June 2021.

[13] Francesco Croce and Matthias Hein. Minimally distorted adversarial examples with a fast adaptive boundary attack. In *International Conference on Machine Learning*, 2020. https://arxiv.org/abs/1907.02044.

[14] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International Conference on Machine Learning*, 2020. https://arxiv.org/abs/2003.01690.

[15] Edward E. Cureton. The normal approximation to the signed-rank sampling distribution when zero differences are present. In *Journal of the American Statistical Association*, volume 62, pages 1068–1069, 1967. https://www.tandfonline.com/doi/abs/10.1080/01621459.1967.10500917.

[16] Gavin Weiguang Ding, Yash Sharma, Kry Yik Chau Lui, and Ruitong Huang. Mma training: Direct input space margin maximization through adversarial training. In *International Conference on Learning Representations*, 2020. https://openreview.net/forum?id=HkeryxBtPB.

[17] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Conference on Computer Vision and Pattern Recognition*, 2018. https://arxiv.org/abs/1710.06081.

[18] Reuben Feinman, Ryan R Curtin, Saurabh Shintre, and Andrew B Gardner. Detecting adversarial samples from artifacts. *arXiv preprint arXiv:1703.00410*, 2017.

[19] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015. https://arxiv.org/abs/1412.6572.

[20] Sven Gowal, Chongli Qin, Jonathan Uesato, Timothy Mann, and Pushmeet Kohli. Uncovering the limits of adversarial training against norm-bounded adversarial examples. *arXiv preprint arXiv:2010.03593*, 2020.

[21] Chuan Guo, Mayank Rana, Moustapha Cisse, and Laurens Van Der Maaten. Countering adversarial images using input transformations. In *International Conference on Learning Representations*, 2018.

[22] Warren He, James Wei, Xinyun Chen, Nicholas Carlini, and Dawn Song. Adversarial example defense: Ensembles of weak defenses are not strong. In *11th {USENIX} workshop on offensive technologies ({WOOT} 17)*, 2017.

[23] Dan Hendrycks, Kimin Lee, and Mantas Mazeika. Using pre-training can improve model robustness and uncertainty. In *International Conference on Machine Learning*, 2019. https://arxiv.org/abs/1901.09960.

[24] Guy Katz, Derek A Huang, Duligur Ibeling, Kyle Julian, Christopher Lazarus, Rachel Lim, Parth Shah, Shantanu Thakoor, Haoze Wu, Aleksandar Zeljić, et al. The marabou framework for verification and analysis of deep neural networks. In *International Conference on Computer Aided Verification*, 2019.

[25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. https://arxiv.org/abs/1412.6980.

[26] Aounon Kumar, Alexander Levine, Tom Goldstein, and Soheil Feizi. Curse of dimensionality on randomized smoothing for certifiable robustness. In *International Conference on Machine Learning*, 2020.

[27] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *International Conference on Learning Representations Workshop*, 2017. https://arxiv.org/abs/1607.02533.

[28] Alexey Kurakin, Ian Goodfellow, Samy Bengio, Yinpeng Dong, Fangzhou Liao, Ming Liang, Tianyu Pang, Jun Zhu, Xiaolin Hu, Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, Alan Yuille, Sangxia Huang, Yao Zhao, Yuzhe Zhao, Zhonglin Han, Junjiajia Long, Yerkebulan Berdibekov, Takuya Akiba, Seiya Tokui, and Motoki Abe. Adversarial attacks and defences competition. In *Conference on Neural Information Processing Systems*, 2017. https://arxiv.org/abs/1804.00097.

[29] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial machine learning at scale. In *International Conference on Learning Representations*, 2017. https://arxiv.org/abs/1611.01236.

[30] Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy. In *2019 IEEE Symposium on Security and Privacy (SP)*, 2019.

[31] Fangzhou Liao, Ming Liang, Yinpeng Dong, Tianyu Pang, Jun Zhu, and Xiaolin Hu. Defense against adversarial attacks using high-level representation guided denoiser. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1778–1787, 2018.

[32] Xingjun Ma, Bo Li, Yisen Wang, Sarah M Erfani, Sudanthi Wijewickrema, Grant Schoenebeck, Dawn Song, Michael E Houle, and James Bailey. Characterizing adversarial subspaces using local intrinsic dimensionality. In *International Conference on Learning Representations*, 2018.

[33] 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*, 2018. https://arxiv.org/pdf/1706.06083.

[34] Jan Hendrik Metzen, Tim Genewein, Volker Fischer, and Bastian Bischoff. On detecting adversarial perturbations. In *International Conference on Learning Representations*, 2017. https://arxiv.org/abs/1702.04267.

[35] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Conference on Computer Vision and Pattern Recognition*, 2016. https://arxiv.org/abs/1511.04599.

[36] Seyed-Mohsen Moosavi-Dezfooli, Ashish Shrivastava, and Oncel Tuzel. Divide, denoise, and defend against adversarial attacks. In *Conference on Computer Vision and Pattern Recognition*, 2018. https://arxiv.org/abs/1802.06806.

[37] Marius Mosbach, Maksym Andriushchenko, Thomas Alexander Trost, Matthias Hein, and Dietrich Klakow. Logit pairing methods can fool gradient-based attacks. In *NeurIPS 2018 Workshop on Security in Machine Learning*, 2018. https://arxiv.org/abs/1810.12042.

[38] Nicolas Papernot, Patrick Mcdaniel, Ian J. Goodfellow, Somesh Jha, Z. Berkay Celik, and Ananthram Swami. Practical black-box attacks against deep learning systems using adversarial examples. In *ACM ASIA Conference on Computer and Communications Security*, 2017. https://arxiv.org/abs/1602.02697.

[39] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z. Berkay Celik, and Ananthram Swami. The limitations of deep learning in adversarial settings. In *2016 IEEE European Symposium on Security and Privacy (EuroS P)*, pages 372–387, 2016.

[40] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. Certified defenses against adversarial examples. In *International Conference on Learning Representations*, 2018. https://arxiv.org/abs/1801.09344.

[41] Sylvestre-Alvise Rebuffi, Sven Gowal, Dan A Calian, Florian Stimberg, Olivia Wiles, and Timothy Mann. Fixing data augmentation to improve adversarial robustness. *arXiv preprint arXiv:2103.01946*, 2021.

[42] Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do adversarially robust imagenet models transfer better? In *Conference on Neural Information Processing Systems*, 2020. https://arxiv.org/abs/2007.08489.

[43] Pouya Samangouei, Maya Kabkab, and Rama Chellappa. Defense-GAN: Protecting classifiers against adversarial attacks using generative models. In *International Conference on Learning Representations*, 2018.

[44] Vikash Sehwag, Shiqi Wang, Prateek Mittal, and Suman Jana. Hydra: Pruning adversarially robust neural networks. In *Conference on Neural Information Processing Systems*, 2020. https://arxiv.org/abs/2002.10509.

[45] Ayon Sen, Xiaojin Zhu, Liam Marshall, and Robert Nowak. Should adversarial attacks use pixel p-norm? arXiv preprint 1906.02439, June 2019.

[46] Ali Shafahi, Mahyar Najibi, Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S. Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! In *Conference on Neural Information Processing Systems*, 2019. https://arxiv.org/abs/1904.12843.

[47] Mahmood Sharif, Lujo Bauer, and Michael K. Reiter. On the suitability of lp-norms for creating and preventing adversarial examples. In *Conference on Computer Vision and Pattern Recognition Workshops*, 2018. https://arxiv.org/abs/1802.09653.

[48] Gagandeep Singh, Timon Gehr, Markus Püschel, and Martin Vechev. An abstract domain for certifying neural networks. *Proceedings of the ACM on Programming Languages*, 3(POPL):1–30, 2019.

[49] Yang Song, Taesup Kim, Sebastian Nowozin, Stefano Ermon, and Nate Kushman. Pixeldefend: Leveraging generative models to understand and defend against adversarial examples. In *International Conference on Learning Representations*, 2018. https://arxiv.org/abs/1710.10766.

[50] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014. https://arxiv.org/abs/1312.6199.

[51] Yusuke Tashiro, Yang Song, and Stefano Ermon. Diversity can be transferred: Output diversification for white-and black-box attacks. In *Conference on Neural Information Processing Systems*, 2020. .

[52] Florian Tramèr, Jens Behrmann, Nicholas Carlini, Nicolas Papernot, and Jörn-Henrik Jacobsen. Fundamental tradeoffs between invariance and sensitivity to adversarial perturbations. In *International Conference on Machine Learning*, 2020. https://arxiv.org/abs/2002.04599.

[53] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. In *International Conference on Learning Representations*, 2018. https://arxiv.org/abs/1705.07204.

[54] Jonathan Uesato, Brendan O'Donoghue, Aäron van den Oord, and Pushmeet Kohli. Adversarial risk and the dangers of evaluating against weak attacks. *International Conference on Machine Learning*, 2018. https://arxiv.org/abs/1802.05666.

[55] Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. Improving adversarial robustness requires revisiting misclassified examples. In *International Conference on Learning Representations*, 2020. https://openreview.net/forum?id=rklOg6EFwS.

[56] Eric Wong and Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning*, 2018.

[57] Eric Wong, Leslie Rice, and J. Zico Kolter. Fast is better than free: Revisiting adversarial training. In *International Conference on Learning Representations*, 2020. https://arxiv.org/abs/2001.03994.

[58] Dongxian Wu, Shu tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. In *Conference on Neural Information Processing Systems*, 2020. https://arxiv.org/abs/2004.05884.

[59] Chaowei Xiao, Bo Li, Jun-Yan Zhu, Warren He, M. Liu, and Dawn Xiaodong Song. Generating adversarial examples with adversarial networks, 2018.

[60] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. In *The Network and Distributed System Security Symposium (NDSS)*, 2018. https://arxiv.org/abs/1704.01155.

[61] Huan Zhang, Hongge Chen, Chaowei Xiao, Bo Li, Duane S. Boning, and Cho-Jui Hsieh. Towards stable and efficient training of verifiably robust neural networks. In *International Conference on Learning Representations*, 2020. https://arxiv.org/abs/1906.06316.

A Performance

As we demonstrated in Sec. 6.1, on average, Auto-PGD with MD loss finds more adversarial examples than Auto-PGD with all three other loss functions, while Auto-PGD with MD loss and smart rounding, performs slightly better than without

smart rounding, and CGD attacks performs better than Auto-PGD with MD loss and smart rounding. We observe similar results across datasets, values of ε, and defenses, as we show in Figs. 7–9.

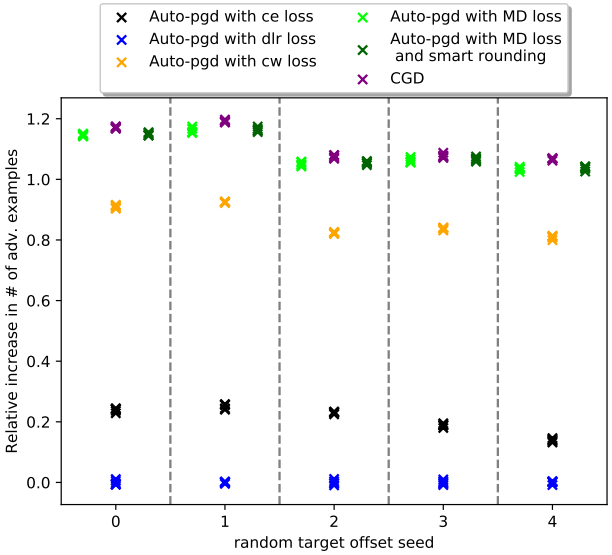


Figure 7: This figure shows the relative improvement in the number of adversarial examples found by attacks on different defenses compared to the worst-performing attack. Experiments were performed using 50,000 images from the testing set of ImageNet. We ran attacks using $\varepsilon = 8/255$ against SIE+20 [42], five different random initial perturbations with seeds 0–4, and five different random target offsets with seeds 0–4. The result is normalized by the mean of the worst performing method of each target offset.

B Uniqueness of Attack Methods

Each attack method finds some unique adversarial examples compared to others. The only exception is between Auto-PGD using MD loss, with and without smart rounding, as the difference is additional tests in each iteration. The adversarial examples that Auto-PGD with MD loss but without smart rounding finds comprise a subset, although not necessarily a proper subset, of the adversarial examples that Auto-PGD with MD loss and smart rounding finds. We observe the same phenomena across attack methods, values of ε, defenses, and datasets, as shown in Figs. 10–26.

C Effects of Random Initial Perturbations

As described in Sec. 6.3, each attack method finds more adversarial examples when they are allowed to use more different initial perturbations, up to the 20 that we have tried per image.

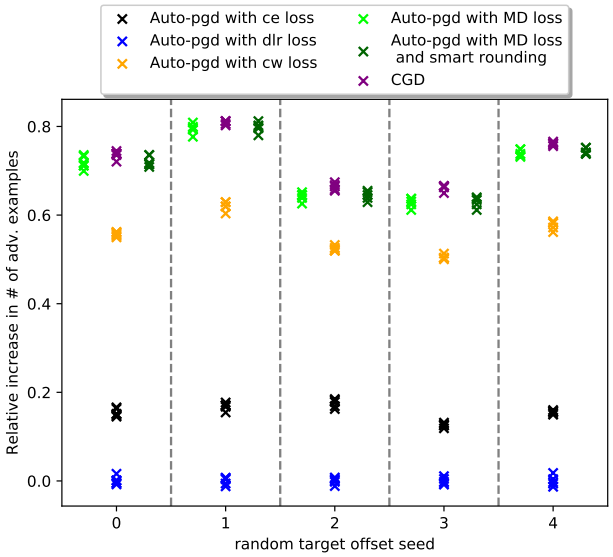


Figure 8: This figure shows the relative improvement in the number of adversarial examples found by attacks on different defenses compared to the worst-performing attack. Experiments were performed using 50,000 images from the testing set of ImageNet. We ran attacks using $\varepsilon = 4/255$ against SIE+20 [42], five different random initial perturbations with seeds 0–4, and five different random target offsets with seeds 0–4. The result is normalized by the mean of the worst performing method of each target offset.

We observe the same phenomena across attack methods, values of ε, and defenses on CIFAR10, as shown in Figs. 27–39.

D Untargeted Attacks

As we introduced in Sec. 7.2, on average, $\text{CGD}_{\text{untarg}}$ performed better than Auto-PGD with L_{DLR} and Auto-PGD with L_{CW^*}, each with or without smart rounding. Figs. 40–41 show similar results when we use different values of ε.

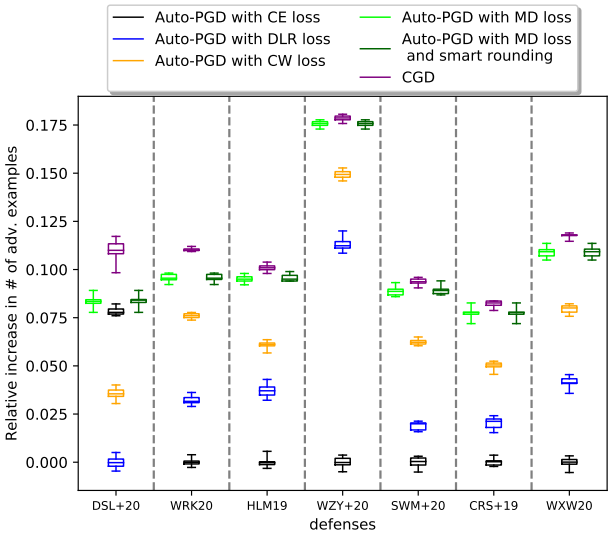


Figure 9: This figure shows the relative improvement in the number of adversarial examples found by attacks on different defenses compared to the worst-performing attack. Experiments were performed using 10,000 images from the testing set of CIFAR10. We ran attacks using $\varepsilon = 8/255$, 20 different random initial perturbations with seeds 0–19, and a fixed random target offset with seed 0. The result is normalized by the mean of the worst performing method against each model.

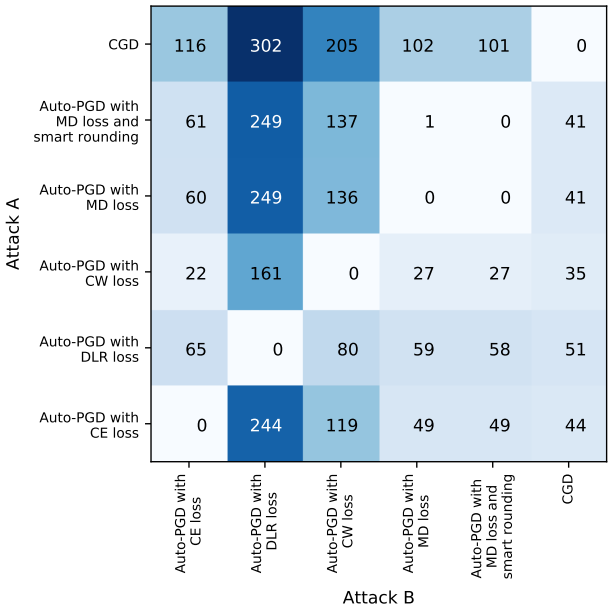


Figure 10: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 8/255$ against the DSL+20 [16] defense, no decimals are kept.

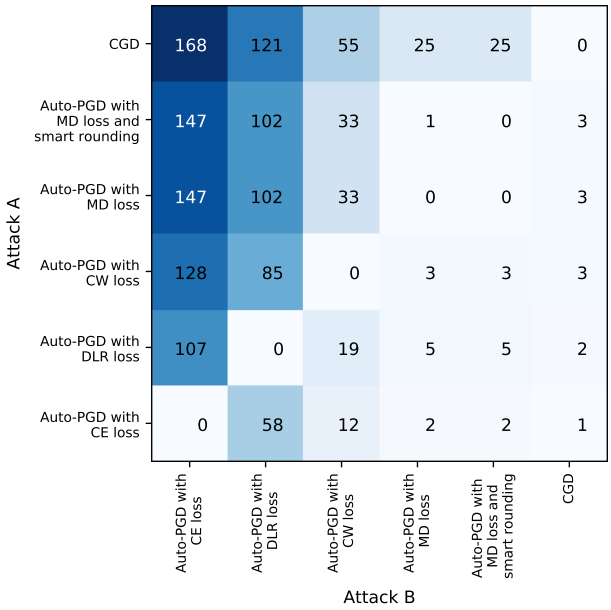


Figure 11: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 8/255$ against the WRK20 [57] defense, no decimals are kept.

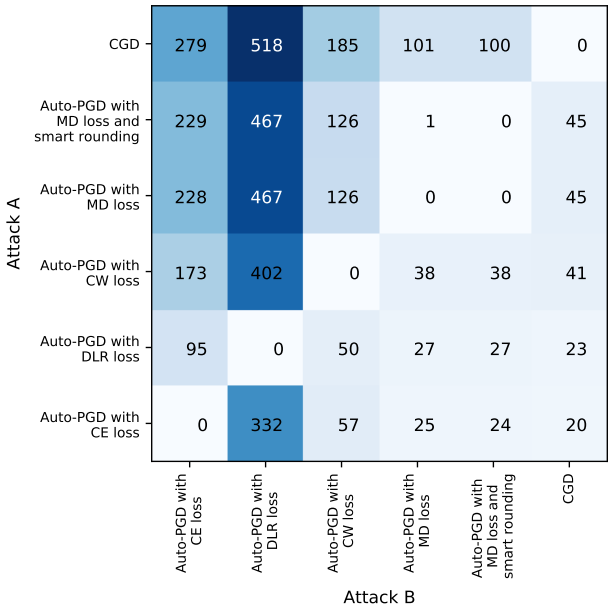


Figure 12: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 16/255$ against the WRK20 [57] defense, no decimals are kept

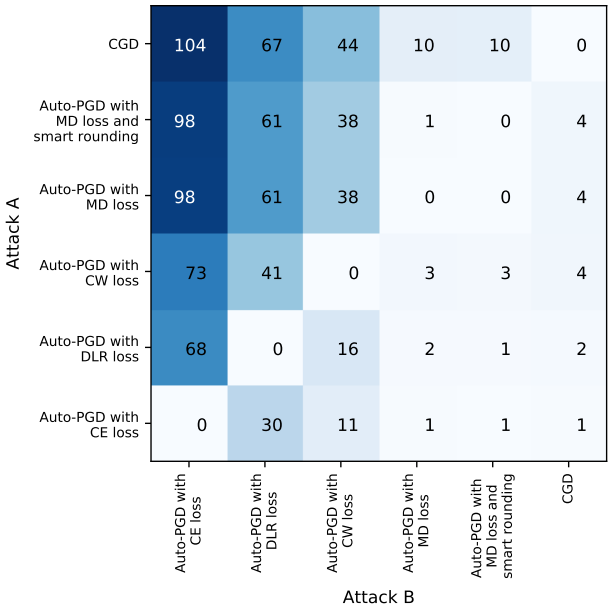


Figure 13: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 8/255$ against the HLM19 [23] defense, no decimals are kept.

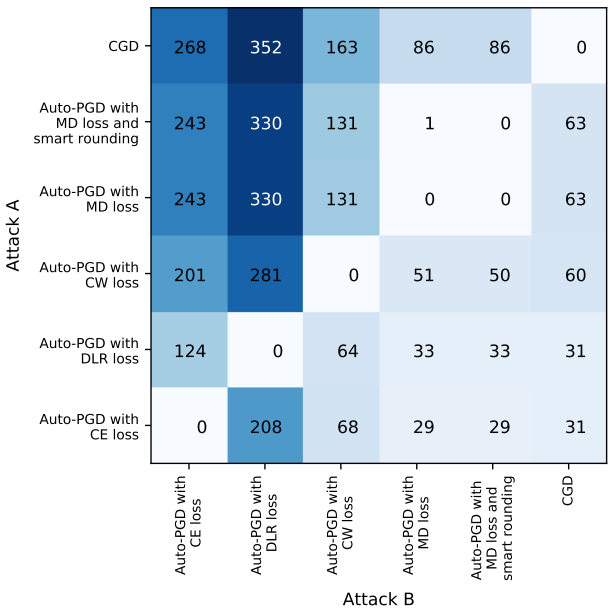


Figure 14: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 16/255$ against the HLM19 [23] defense, no decimals are kept.

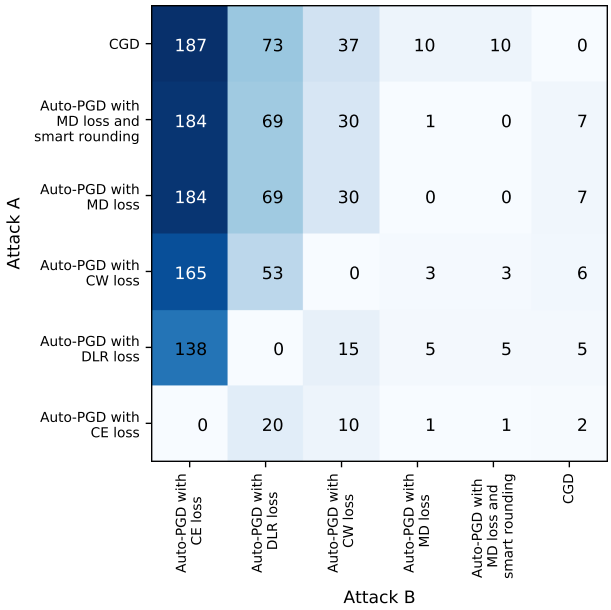


Figure 15: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 8/255$ against the WZY+20 [55] defense, no decimals are kept.

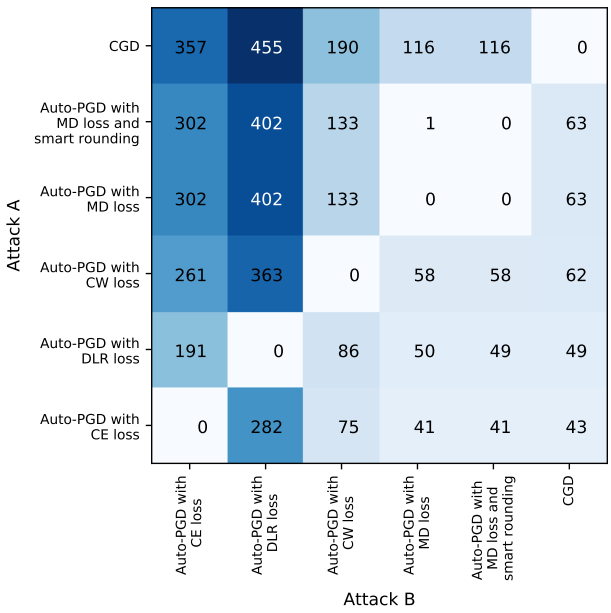


Figure 16: This image shows the average number of successful adversarial examples found by attack A but not by attack B each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 16/255$ against the WZY+20 [55] defense, no decimals are kept.

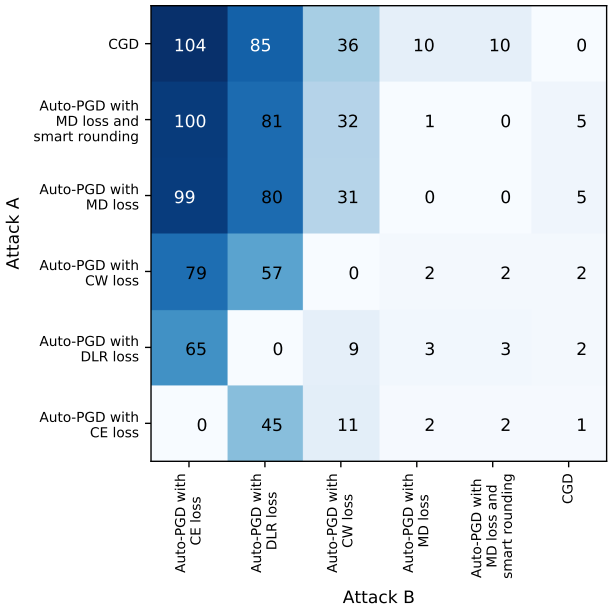


Figure 17: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 8/255$ against the SWM+20 [44] defense, no decimals are kept.

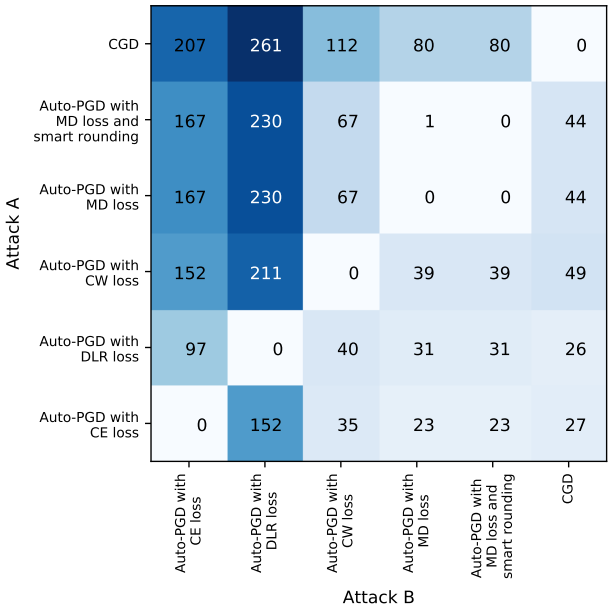


Figure 18: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 16/255$ against the SWM+20 [44] defense, no decimals are kept.

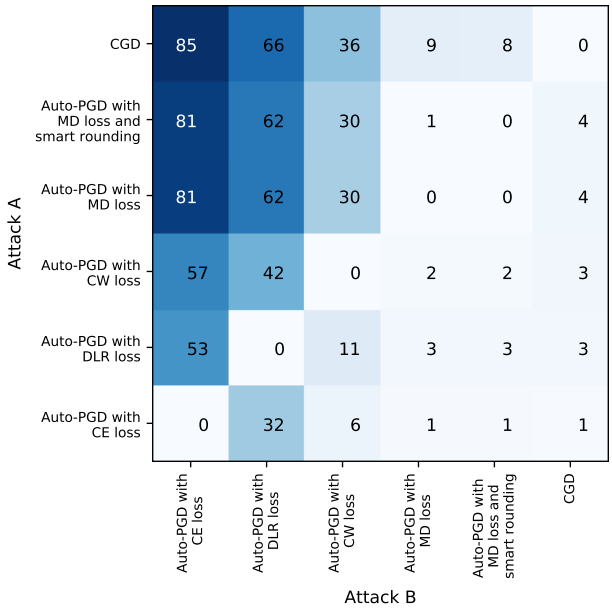


Figure 19: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 8/255$ against the CRS+19 [9] defense, no decimals are kept.

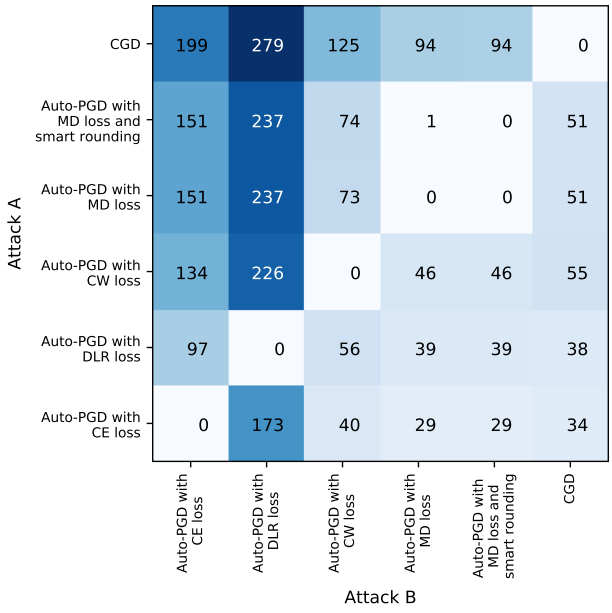


Figure 20: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 16/255$ against the CRS+19 [9] defense, no decimals are kept.

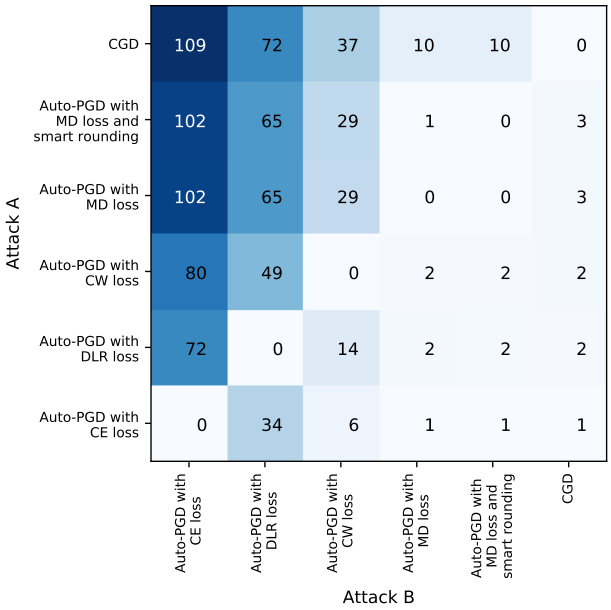


Figure 21: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 8/255$ against the WXW20 [58] defense, no decimals are kept.

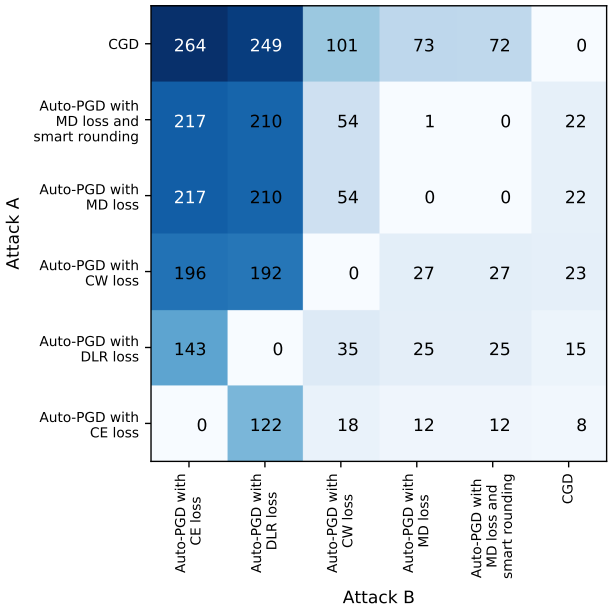


Figure 22: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 16/255$ against the WXW20 [58] defense, no decimals are kept.

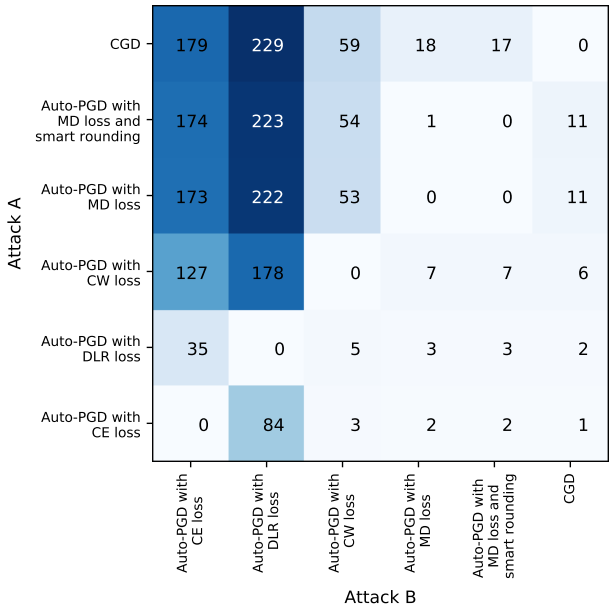


Figure 23: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 50,000 attempts on the testing set of ImageNet using $\varepsilon = 4/255$ against the SIE+20 [42] defense, no decimals are kept.

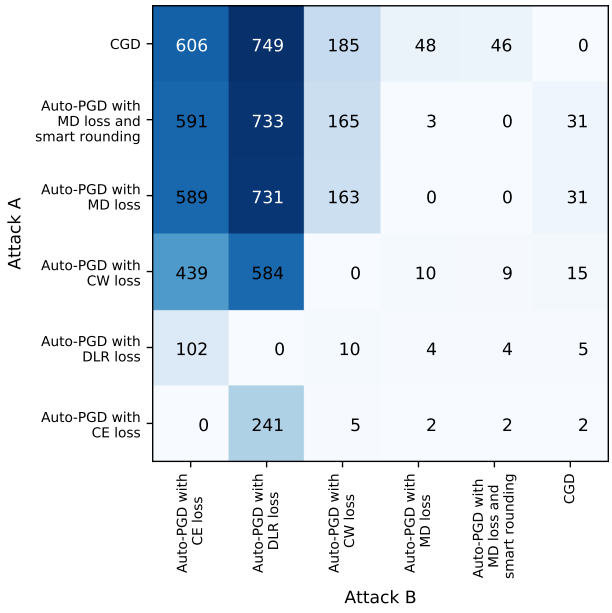


Figure 24: This image shows the average number of successful adversarial examples found by attack A but not by attack B in each set of 50,000 attempts on the testing set of ImageNet using $\varepsilon = 8/255$ against the SIE+20 [42] defense, no decimals are kept.

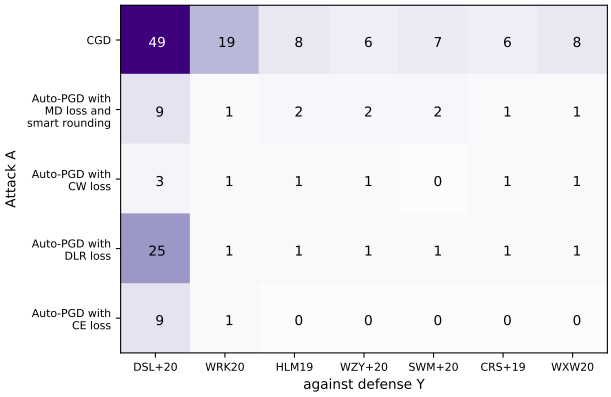


Figure 25: This is the average number of successful adversarial examples found by attack A but not by any of other attack methods, against each defense Y in each set of 10,000 attempts on the testing set of CIFAR10 using $\varepsilon = 8/255$, no decimals are kept. Auto-PGD using MD loss but without smart rounding is not included in this plot as the adversarial examples it found are a subset of those found by Auto-PGD using MD loss and with smart rounding.

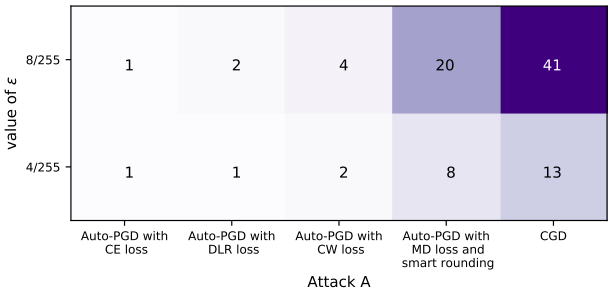


Figure 26: This is the average number of successful adversarial examples found by attack A but not by any other attack methods, against the SIE+20 [42] defense in each set of 50,000 attempts on the testing set of ImageNet, no decimals are kept. Auto-PGD using MD loss but without smart rounding is not included in this plot as the adversarial examples it found are a subset of those found by Auto-PGD using MD loss and with smart rounding.

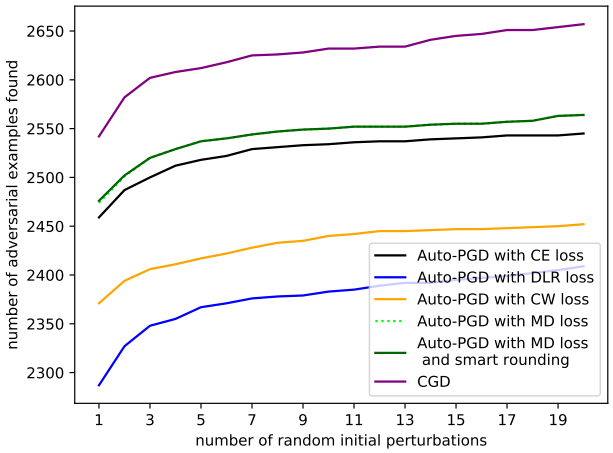


Figure 27: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 8/255$, against the DSL+20 [16]defense.

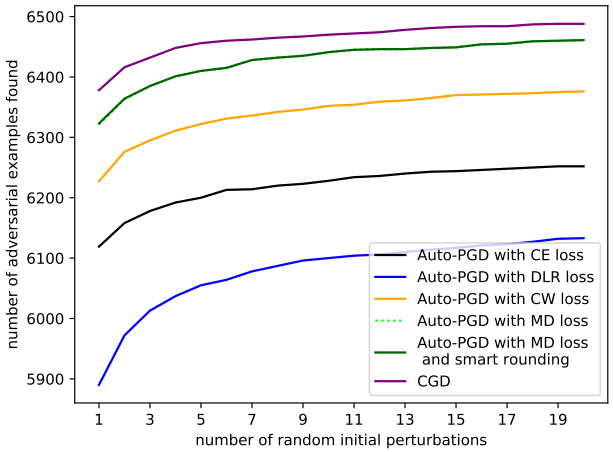


Figure 28: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 16/255$, against the WRK20 [57] defense.

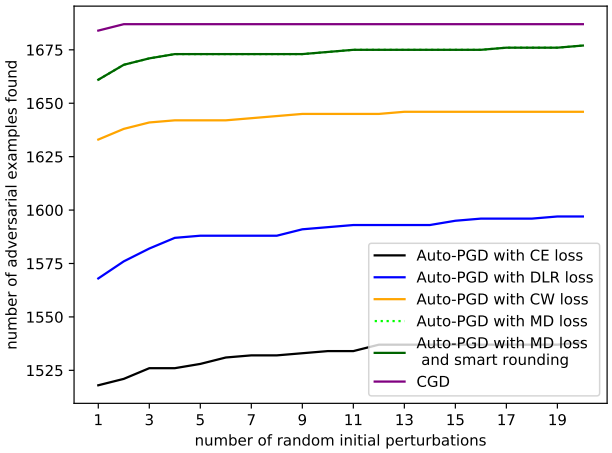


Figure 29: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 8/255$, against the WRK20 [57] defense.

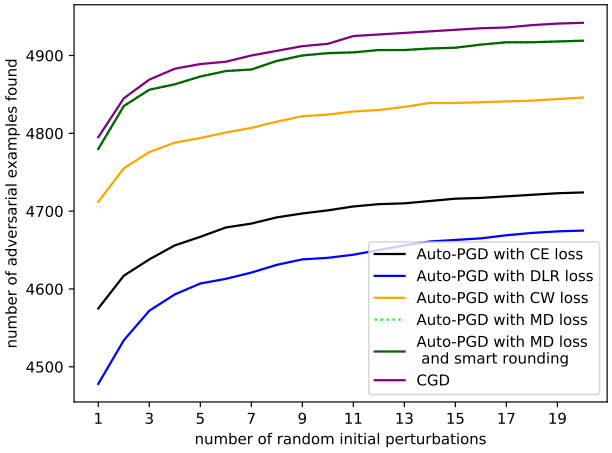


Figure 30: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 16/255$, against the HLM19 [23] defense.

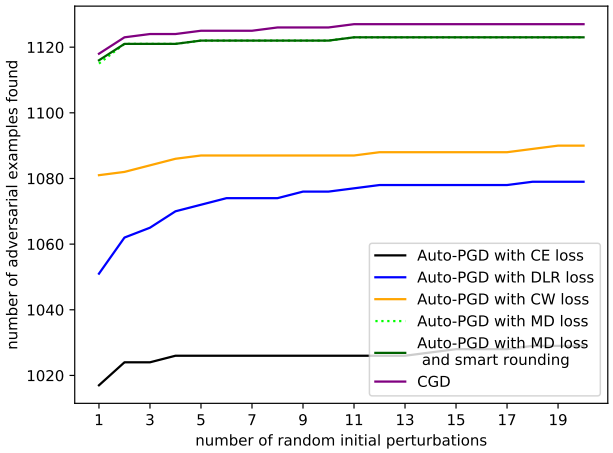


Figure 31: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 8/255$, against the HLM19 [23] defense.

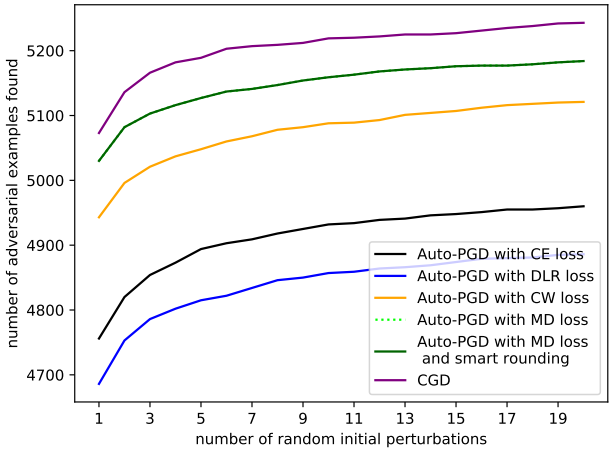


Figure 32: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 16/255$, against the WZY+20 [55] defense.

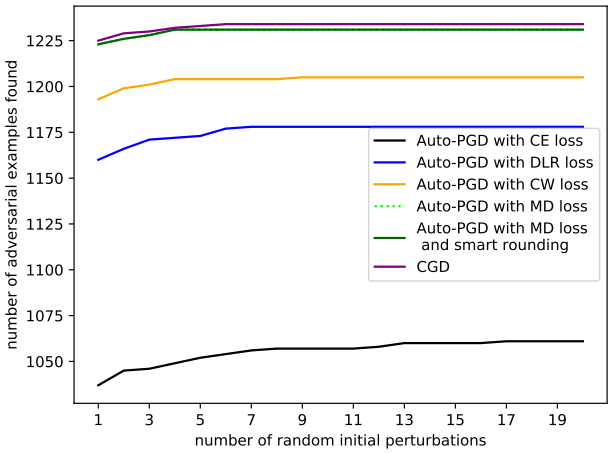


Figure 33: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 8/255$, against the WZY+20 [55] defense.

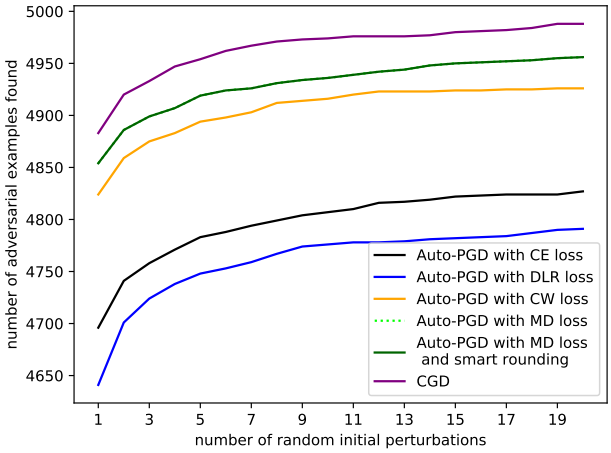


Figure 34: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 16/255$, against the SWM+20 [44] defense.

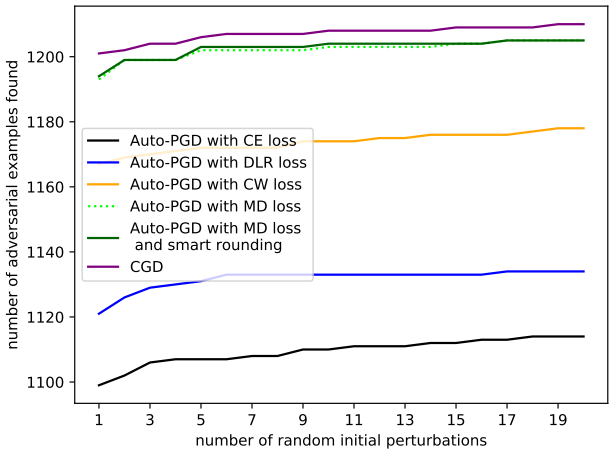


Figure 35: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 8/255$, against the SWM+20 [44] defense.

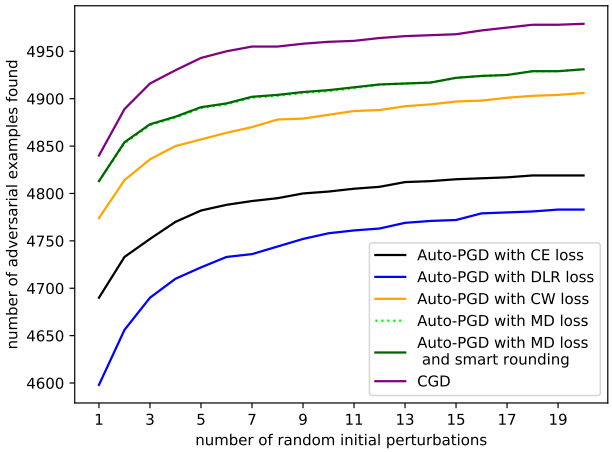


Figure 36: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 16/255$, against the CRS+19 [9]defense.

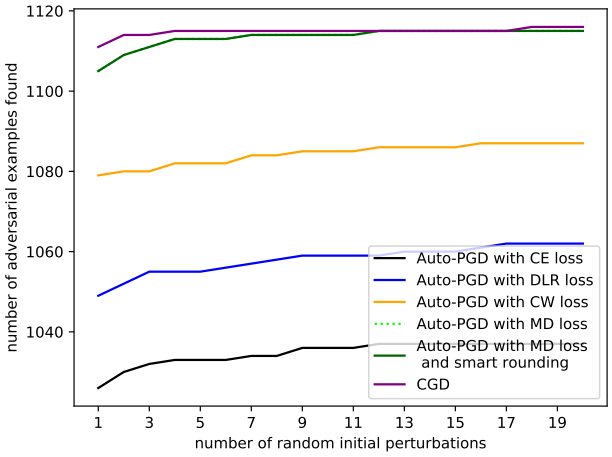


Figure 37: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 8/255$, against the CRS+19 [9] defense.

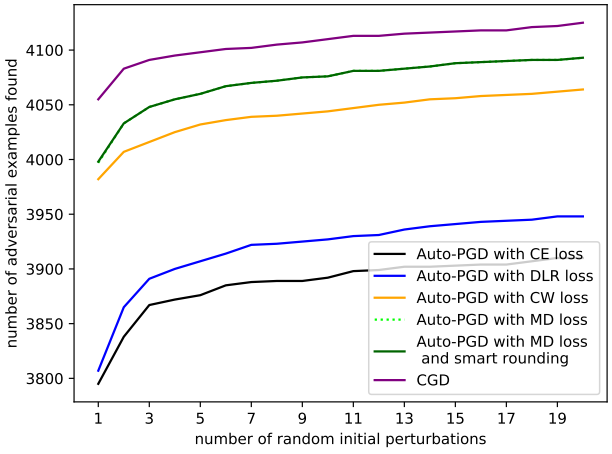


Figure 38: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 16/255$, against the WXW20 [58] defense.

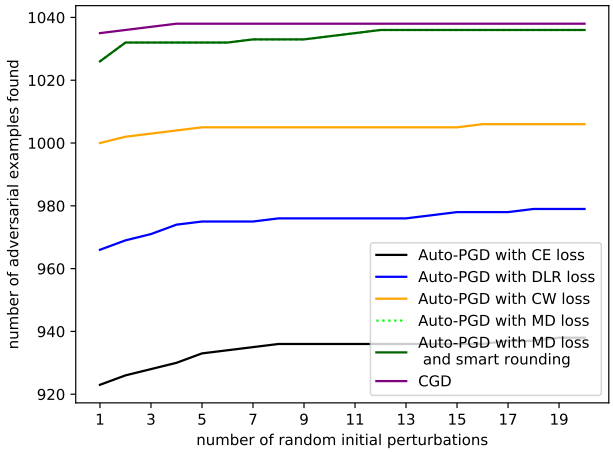


Figure 39: The image shows the number of adversarial examples each of the attacks found when they are allowed to use the specified number of random initial perturbations, on all 10,000 images from the testing set of CIFAR10, with $\varepsilon = 8/255$, against the WXW20 [58] defense.

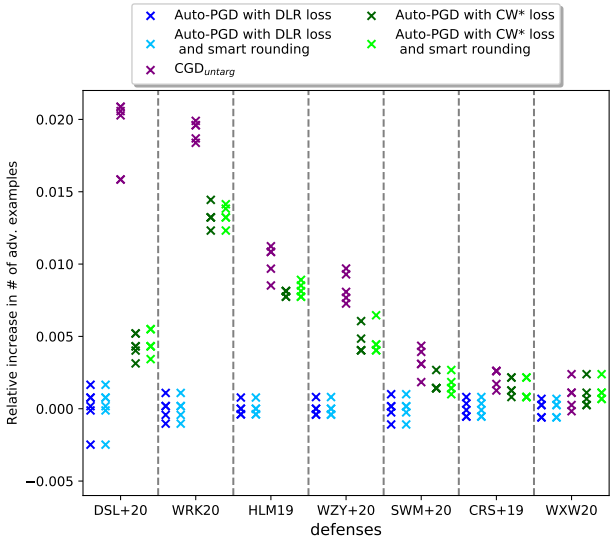


Figure 40: This figure shows the relative improvement in the number of adversarial examples found by untargeted attacks on different defenses compared to the worst-performing attack. Experiments were performed using 10,000 images from the testing set of CIFAR10. We ran attacks using $\varepsilon = 4/255$, and five different random initial perturbations with seeds 0–4. The result is normalized by the mean of the worst performing method against each model.

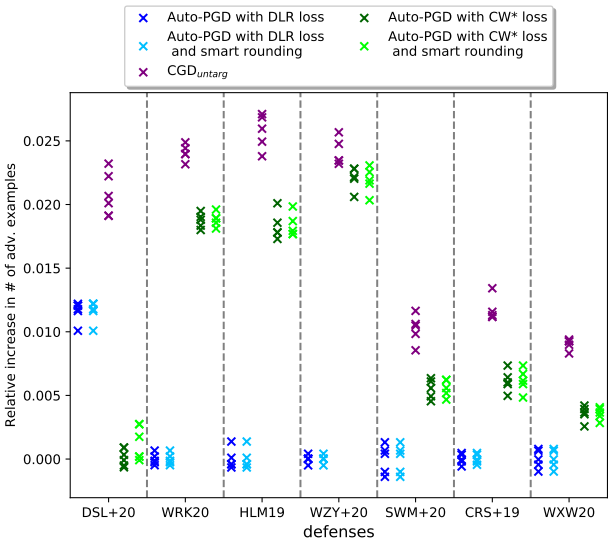


Figure 41: This figure shows the relative improvement in the number of adversarial examples found by untargeted attacks on different defenses compared to the worst-performing attack. Experiments were performed using 10,000 images from the testing set of CIFAR10. We ran attacks using $\varepsilon = 16/255$, and five different random initial perturbations with seeds 0–4. The result is normalized by the mean of the worst performing method against each model.