Alternating Objectives Generates Stronger PGD-Based Adversarial Attacks

Antoniou Nikolaos [*], Efthymios Georgiou [*†], Alexandros Potamianos [*]

[*] School of Electrical and Computer Engineering, National Technical University of Athens, Athens, Greece
[†] Institute for Language and Speech Processing, Athena Research Center, Athens, Greece
antoniou_nikos@hotmail.com, efthygeo@mail.ntua.gr, potam@central.ntua.gr

Abstract—Designing powerful adversarial attacks is of paramount importance for the evaluation of ℓ_p-bounded adversarial defenses. Projected Gradient Descent (PGD) is one of the most effective and conceptually simple algorithms to generate such adversaries. The search space of PGD is dictated by the steepest ascent directions of an objective. Despite the plethora of objective function choices, there is no universally superior option and robustness overestimation may arise from ill-suited objective selection. Driven by this observation, we postulate that the combination of different objectives through a simple loss alternating scheme renders PGD more robust towards design choices. We experimentally verify this assertion on a synthetic-data example and by evaluating our proposed method across 25 different ℓ_∞-robust models and 3 datasets. The performance improvement is consistent, when compared to the single loss counterparts. In the CIFAR-10 dataset, our strongest adversarial attack outperforms all of the white-box components of AutoAttack (AA) ensemble [1], as well as the most powerful attacks existing on the literature, achieving state-of-the-art results in the computational budget of our study ($T = 100$, no restarts).

Index Terms—Adversarial Attacks, ℓ_∞-bounded robustness, Projected Gradient Descent, RobustBench Benchmark

I. INTRODUCTION

The advent of Deep Learning (DL) caused a paradigm shift and revolutionized the way that various interesting applications are approached. Such a wide adoption, however, demands from the research community to comprehend the scenarios where Deep Neural Networks (DNNs) malfunction. This necessity becomes even more imperative when considering the abundance of safety-critical applications that do not leave room for complacency, e.g., autonomous driving. Unfortunately, DNNs have significant failure modes and behave counterintuitively. A prominent instance of this behaviour is illustrated by Szegedy et al. [2], where they showcase that DNN-based image classifiers are vulnerable against *adversarial examples*. These examples arise from applying humanly imperceptible perturbations to clean images, which are capable of degrading the model's predictive performance. This finding triggered research interest on two fronts: *Adversarial Attacks*, which are algorithms to generate such malicious examples and *Adversarial Defenses*, which are methods of increasing the robustness of neural networks. Adversarial robustness is primarily studied through the ℓ_p-bounded threat model, where the perturbation's ℓ_p-norm is bounded by a small constant.

The robustness of Adversarial Defenses, on a given dataset, is estimated by the rate of test set's adversarial examples that the defense can properly classify. Of course, the estimated rate (also called robust accuracy) depends on the strength of the attacking algorithm that will be used for evaluation. Employing weak attacks to evaluate robustness creates a false sense of security, an issue widely known as robustness overestimation [3]–[5].

Arguably, Projected Gradient Descent (PGD) is the most popular adversarial attack used to evaluate ℓ_p−bounded robustness. PGD operates by iteratively following the steepest ascent directions of an objective function, often called the surrogate. PGD has raised in many guises in the adversarial attack literature: Goodfellow et al. [6] propose to attack networks through the Fast Gradient Sign Method (FGSM), which takes a single normalized step, i.e., applying the sign function in the case of ℓ_∞-norm, towards the steepest ascent direction. Kurakin et al. [7] demonstrate that the multi-step variants of FGSM are capable of producing significantly stronger attacks. Dong et al. [8] suggest a modification of the iterative FGSM that integrates a momentum term. Madry et al. [9] link the iterative FGSM with the classical optimization algorithm of PGD.

Despite that PGD combines both simplicity (in terms of implementation) and strength, it has been shown that its performance can be hindered by ill-suited selection of hyperparameters, e.g., fixed step size [1]. Another hyperparameter of consideration is the surrogate loss, for which literature has converged into 3 options: Cross-Entropy (CE) [6], [9], Margin (a.k.a. CW) loss [10] and the Difference of Logits Ratio (DLR) loss [1], with the appealing property of scale-invariance. However, empirical evidence (e.g., as in Figures 9-11 of [1]) shows that there is no universally superior objective and its effectiveness depends on the architecture, weights, training dataset etc. On top of this, certain choices may be improper in special problematic cases: 1) CE yields zero gradients for inputs where the classifier assigns the entire probability mass to the ground truth class [1], [10], 2) both CE and CW are not scale-invariant hence logit rescalings may induce gradient masking [1] and 3) Ma et al. [11] assert that objectives which involve multiple logit terms, i.e., all three of CE,CW and DLR, may suffer from the problem of *gradient imbalance* where logits have quite disparate magnitudes and one term alone steers the optimization trajectory towards non-

optimal solutions.

In this work, PGD is studied from the perspective of surrogate loss. In order to alleviate potentially weak PGD performance arising from poor surrogate selection, we propose to combine different objectives in the same run of PGD. Hopefully, this combination will render PGD less dependent to the surrogate hyperparameter. We identify that a simple alternation of objectives during PGD is sufficient to induce significant boost on the PGD performance over the single loss variants. Further qualitative analysis implies that the switching between different objectives helps the algorithm to expand its search space, visiting more distant intermediate points during its execution.

In this paper, we make the following key contributions:

- We propose to combine multiple objectives during PGD through alternating between them during optimization, in order to alleviate potential flaws of each objective. Our proposed strategy outperforms the single-loss PGD variants in 25 out of 25 (15 on CIFAR-10, 6 on CIFAR-100 and 4 on ImageNet) tested $\ell_\infty-$bounded robust models.
- For the CIFAR-10 dataset, our attack outperforms the three white-box components of AutoAttack [1]: APGD$_{\text{CE}}$, APGD$_{\text{DLR}}$ and FAB attack [12]. Furthermore, in most cases our attack achieves higher Attack Success Rate (ASR) than the strongest baselines (for $T = 100$ iterations and $R = 1$ restart) in the literature: GAMA-PGD [13] and MD attack [11].
- We present extensive experimentation and analysis regarding the proposed alternation scheme, including: 1) A synthetic example which highlights how PGD with a single loss can fail, 2) Qualitative analysis indicating that switching losses promotes search diversity and 3) Ablation experiments which demonstrate that this loss combination strategy is more effective than two other combining methods.

The remainder of this paper is organized as follows: Section II provides the necessary background, covering basic aspects of the worst-case ℓ_p-bounded adversarial robustness, Section III briefly discusses research work related to PGD-based attacks, since PGD is the main topic of our study. In Section V we conduct numerous experiments to verify the effectiveness of our proposed method, whereas in Section VI we discuss how our study differs from previous related work.

II. BACKGROUND

A. Notation

Image-label pairs are denoted as $(\mathbf{x}, y) \in \mathcal{X} \times \mathcal{Y}$ where $\mathcal{X} \subseteq \mathbb{R}^D, \mathcal{Y} \subseteq \mathbb{Z}$. The classifier's logit representation will be denoted as $\mathbf{z}(\mathbf{x}) \in \mathbb{R}^C$ (or simply $\mathbf{z}$), where C: the total number of classes. Applying a softmax layer to the logit vector produces the probability vector $p(y|\mathbf{x})$. The classification decision will be denoted as $f(\mathbf{x})$, hence $f(\mathbf{x}) = \arg\max_{i \in [C]} \mathbf{z}(\mathbf{x})_i$, where $[C] = \{1, ..., C\}$. The surrogate loss $\mathcal{L}(\mathbf{z}(\mathbf{x}), y)$ (which will also be referred as $\mathcal{L}(\mathbf{x}, y)$, for brevity's sake), e.g., cross-entropy, measures the model's ability to assign the label y to example $\mathbf{x}$.

B. Threat Model

The constraint of creating an imperceptible perturbation is approximated through the bounded ℓ_p-norm condition. The generation of adversarial attacks should obey this restriction, returning an output that lies within the ℓ_p-ball of radius ϵ around the clean input $\mathbf{x}$. Hence, the search space of potential adversaries for the image $\mathbf{x}$ can be expressed as:

$$\mathcal{S}(\mathbf{x}) = \{\mathbf{x}' : \|\mathbf{x} - \mathbf{x}'\|_p \leq \epsilon\} \quad (1)$$

Despite that the ℓ_p-bounded threat model is only a crude approximation of true similarity between data samples like images, solving the problem of ℓ_p-bounded robustness can be viewed as an important stepping stone towards confronting more realistic scenarios.

C. A taxonomy of ℓ_p-bounded adversarial attacks

Next we present a basic categorization of adversarial attacks based on their capabilities during generation and their end goal.

Adversary's Knowledge. Based on the amount of information that the adversary has at its disposal, attacks can be divided into two major categories: *white-box* and *black-box*. In the former, the attacker has access to every aspect of the model: its architecture, weights and training data. This allows the adversary to obtain the network's gradients w.r.t. the input which is particularly useful when creating attacks. In the latter category, however, the adversary can only use the model as an oracle, feeding an input point and getting access to the output vector, or sometimes just to the output class.

Despite that typical real-world scenarios are more similar to the black-box setting, white-box attacks constitute a much more stronger threat model. Therefore, the evaluation of adversarial defenses is typically performed based on white-box attacks.

Low Confidence vs Low Distortion. Attacks are also divided into minimum-confidence and minimum-norm. In the former, the attack algorithm is based on the following formulation, for the input-label pair $(\mathbf{x}, y)$:

$$\boldsymbol{\delta} : \arg\max_{\boldsymbol{\delta}} \mathcal{L}_{0/1}(f(\mathbf{x} + \boldsymbol{\delta}), y) \text{ s.t. } \mathbf{x} + \boldsymbol{\delta} \in \mathcal{S}(\mathbf{x}) \quad (2)$$

where $\mathcal{L}_{0/1}(f(\mathbf{x}), y) = \mathbb{1}[f(\mathbf{x}) \neq y]$ is the 0-1 loss, which due to its discontinuity is replaced by some surrogate loss $\mathcal{L}$ such as cross-entropy. These attacks aim to reduce the ground truth label's confidence as much as possible by spending the entire attack budget ϵ, hence they typically lie on the boundary surface of the feasible set $\mathcal{S}$. The most prominent examples of minimum-confidence adversarial attacks is the Fast Gradient Sign Method (FGSM) [6], the Iterative-FGSM [7] and Projected Gradient Descent (PGD) [9].

Minimum-norm attacks aspire to find the smallest possible perturbation that leads to misclassification:

$$\boldsymbol{\delta} : \arg\min_{\boldsymbol{\delta}} \|\boldsymbol{\delta}\|_p \text{ s.t. } f(\mathbf{x} + \boldsymbol{\delta}) \neq y \qquad (3)$$

where y: the ground-truth label of $\mathbf{x}$. Such attacks usually find adversaries that are within smaller ℓ_p-distance from the clean input $\mathbf{x}$ than the perturbation bound ϵ. Popular examples of this category are: Carlini-Wagner (CW) attack [10], DDN attack [14], Fast Minimum Norm (FMN) [15] and Fast Adaptive Boundary (FAB) attack [12] among others.

Untargeted vs Targeted. Another criterion of dividing adversarial attacks is whether the adversary desires to force a specific label to the attack. In *targeted* attacks, the attack is considered successful if the corresponding adversarial example is classified into a certain target class. In *untargeted* attacks, the goal is simply to produce an example which is incorrectly classified, with no constraint on its new label. Usually, the transition between the two categories is as simple as slightly modifying the objective function, i.e., from descending the target label's confidence to ascending the ground-truth label's confidence.

D. Empirical Adversarial Defenses

Training ℓ_p-robust neural networks, i.e., networks that are resilient against ℓ_p-bounded adversarial attacks, is a complicated problem since we aspire to simultaneously realize two goals. First, the classifier is asked to perform well on unseen examples drawn from the same distribution as the examples used during training. An additional requirement is to find networks that produce smooth predictions, assigning the same label to all data residing inside the ℓ_p-ball of such examples. The most standard way of increasing ℓ_p-bounded robustness is Adversarial Training (AT) [6], [9]; in AT, the defender aims to minimize the *robust expected risk*:

$$\mathcal{R}_{\text{rob}}^f(\boldsymbol{\theta}) = \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}\left[\max_{\boldsymbol{\delta}:\|\boldsymbol{\delta}\|_p\leq\epsilon} \mathbb{1}[f_{\boldsymbol{\theta}}(\mathbf{x}+\boldsymbol{\delta})\neq y]\right] \quad (4)$$

The inner expression coincides with the task of finding the worst-case ℓ_p-bounded adversarial example. Madry et al. [9] confront the problem through the first-order method of PGD. An important barrier of this method is the additional computational overhead. The iterative PGD process renders this method costly in terms of compute, hence a line of research aims to increase robustness using one-step adversaries [16]–[19], in order to restrain the overall training time to similar levels as with standard training. Another important work on adversarial defenses is the TRADES framework, introduced by Zhang et al. [20]. The robust expected risk of Equation 4 can be decomposed as the sum of two individual terms. The first term represents the *classification error*, where the optimization searches parameters that generalize well. The other term, dubbed as *boundary error*, can be considered as exerting a regularizing effect, where it imposes decision "smoothness" between inputs inside the same ℓ_p-ball.

Schmidt et al. [21] provide evidence that adversarially training classifiers may require an increasing amount of data. Following this, many works [22]–[24] explore the use of both pseudo-labeled additional data and elaborate data augmentation techniques.

Robustness Overestimation. Evaluating the true degree of ℓ_p-bounded robustness of empirical methods is intractable, since one needs to calculate the average 0-1 risk on a held-out test set. Typically, the defender deploys a strong attacking algorithm to obtain a lower bound on the true risk. However, this trial-and-error technique can provide misleading results. Failing to select a proper attacking algorithm creates an inaccurate sense of security [3], [4], [25]. Importantly, these works propose numerous indicators that demonstrate whether the evaluation suffers from this issue and guidelines of how to properly evaluate a defense.

The introduction of RobustBench [26], based on the AutoAttack ensemble (comprised of three white-box [1], [12] and one black-box [27] methods), contributed to a consensus regarding the evaluation of ℓ_p-bounded robustness: A newly proposed defense is first "passed" through an AutoAttack evaluation, and then the defender can also perform adaptive attacks [25], based on potential model-specific weaknesses.

Despite the general adoption of AutoAttack as the standard way to perform first-order robustness evaluations, the community is constantly exploring faster and more powerful attack ensembles [28], [29].

III. RELATED WORK

Projected Gradient Descent (PGD) [7], [9] is the most popular minimum-confidence attack. PGD has been the de facto standard for producing ℓ_p-bounded adversarial attacks, especially in the case of $p = \infty$. In short, PGD can be expressed as:

$$\mathbf{x}^{(t+1)} = \mathcal{P}_{\mathcal{S}(\mathbf{x})}\left[\mathbf{x}^{(t)} + \eta^{(t)}\boldsymbol{\delta}^{(t)}\right] \qquad (5)$$

where $\mathbf{x}^{(t)}$: the iterate, $\eta^{(t)}$: step size, $\boldsymbol{\delta}^{(t)}$: update rule of t-th iteration and $\mathcal{P}_{\mathcal{S}}$: the projection operation, which maps the updated iterate into the feasible region $\mathcal{S}$, which in our case is the ℓ_p-ball of radius ϵ around $\mathbf{x}$. Typically, this procedure is repeated multiple times from different random initializations. For a more comprehensive view of how PGD is used to generate adversarial attacks, we refer to the work of Gowal et al. [30], where they present a "holistic" pseudoalgorithm.

In the following discussion we present how one can manipulate the basic building blocks of PGD, namely the optimizer, step size, initialization strategy and surrogate loss, in order to improve its adversarial generation strength.

Optimizer. The optimizer determines the form of the update rule $\boldsymbol{\delta}^{(t)}$. In its simplest version, assuming the surrogate loss $\mathcal{L}(\mathbf{x}, y)$, PGD follows the steepest direction of unit ℓ_p-norm, e.g., the sign of $\nabla_{\mathbf{x}^{(t)}}\mathcal{L}(\mathbf{x}^{(t)}, y)$ in the case of $p = \infty$, or a simple norm-rescaling when $p = 2$. In the C&W attack

[10], the proposed objective is optimized through Adam [31]. The Adam optimizer has also been leveraged in PGD-based works [4], [30]. Dong et al. [8] suggested the incorporation of momentum [32] in the PGD update rule. Subsequently, Croce and Hein [1] proposed the AutoPGD (APGD) variant, wherein the update term is augmented by momentum. Yamamura et al. [33] developed the Auto Conjugate Gradient (ACG) method, which is an elaborate optimizer, adjusting the update rule based on accumulated gradient information from previous steps. ACG is experimentally shown to outperform APGD for a sizable collection of robust models.

Step Size. Another hyperparameter which affects the performance of PGD is the step size $\eta^{(t)}$. In early works, its value is held constant during the entire optimization procedure, e.g., to $\alpha = \epsilon/4$ for ℓ_∞-attacks in CIFAR-10. Croce and Hein [1] conduct large-scale experiments regarding the optimal fixed value, but one immediate corollary is that it greatly depends on the model. Generally, the common trend is to perform some kind of scheduling, where the step size is gradually reduced over time: In [30], [13], the authors apply ten-fold drops at two intermediate timesteps; Ma et al. [11] propose a cosine-annealing scheme, where the step size decays from 2ϵ to 0. In their recent work, Liu et al. [28] adopt a similar decaying strategy. Another interesting way of manipulating this hyperparameter is as in the AutoPGD method [1]; They initially set it to a large value $\alpha = 2\epsilon$, in order to explore the search space sufficiently well. Then, as the optimization proceeds and the iterate gets closer to some local optimum, the need of a more localized search calls for smaller step sizes. Hence, it is halved in specific checkpoints, according to the optimization progress, i.e., based on whether the objective function is reducing or not.

Initialization. Proper initialization plays also a crucial role in the final performance. Typically, the initial point $\mathbf{x}^{(0)}$ can be either set to the clean image $\mathbf{x}$, or alternatively, random noise may be added to the clean image: $\mathbf{x}^{(0)} = \mathbf{x} + \zeta$, where ζ is drawn from some noise distribution. The attack is then repeated multiple times, initialized from different starting points. Tashiro et al. [34] suggest that random initialization may lead to starting points with nearly identical output space representations, hence the attack generates similar results even if executed for many restarts. Output Diversified Initialization (ODI) [34] counteracts this by maximizing the similarity of starting point's logit vector with a random output direction, in the first few PGD iterations. Recently, Liu et al. [28] introduced Adaptive AutoAttack (A^3), the new state-of-the-art attack ensemble. A^3 uses an adaptive initialization strategy, where the starting points are generated by ODI, but instead of following random output space direction, the vector is selected according to prior knowledge of perturbations that led to misclassification.

Surrogate Loss. The maximization of 0-1 loss is intractable for complex function classes as those represented by deep neural networks [35]. It is common to substitute it with a surrogate, differentiable loss which is amenable to optimization methods. A natural candidate is the cross-entropy (CE) objective, which coincides with the negative log-likelihood of the ground truth class. In their seminal work, Carlini and Wagner [10] tested various formulations, obtaining the best performance for the so called margin (or CW) loss. A shared defect in both of these objectives is the lack of scale-invariance, which may be translated in deteriorated performance due to gradient masking. Croce and Hein [1] introduce the Difference of Logits Ration (DLR) loss, which rescales the margin loss to acquire the property of scale-invariance. Most of the literature involves these three options, whose expressions are included below for completeness:

$$
\begin{aligned}
\mathrm{CE}(\mathbf{x}, y) &= -\log p(y|\mathbf{x}) = -\mathbf{z}_y + \log \sum_{j=1}^{C} \exp(\mathbf{z}_j) \\
\mathrm{CW}(\mathbf{x}, y) &= -\mathbf{z}_y + \max_{j \neq y} \mathbf{z}_j \\
\mathrm{DLR}(\mathbf{x}, y) &= -\frac{\mathbf{z}_y + \max_{j \neq y} \mathbf{z}_j}{\mathbf{z}_{\pi_1} - \mathbf{z}_{\pi_3}}
\end{aligned}
\tag{6}
$$

where $\mathbf{z}_\pi$: the logit vector sorted in descending order. Gowal et al. [30] propose the MultiTargeted PGD variant which divides the iteration budget into runs of equal size, where each run optimizes the targeted margin loss, for a different target label per run. Their experiments indicate that the MultiTargeted strategy exploits more judiciously the given computational budget. Sriramanan et al. [13] augment the standard margin loss expression with a regularization term which is set to the MSE between the logit vector of the adversary and its clean counterpart. The weighting coefficient of MSE term is gradually decayed to zero. Ma et al. [11], in an effort to address the issue of imbalanced gradients, optimize only one of the two margin loss terms for the first half of iterations before switching to the typical expression which contains both terms. In the next restart, they repeat the process by using the other term for the first stage of optimization.

IV. METHODOLOGY

Our work is motivated by the observation that a single surrogate loss is unable to perform equally well across different robust models. Croce and Hein [1] provide strong empirical evidence to back up this argument. Specifically, in their study they investigate the effectiveness of three objectives: CE, CW and DLR. These three aforementioned objectives have expressions that are distinguished by small differences, yet each option can profoundly influence the Attack Success Rate (ASR) of PGD. Of course, this phenomenon is not surprising at all: the optimization space coincides with the high-dimensional pixel space of natural images, hence even just a rescaling that links the CW with DLR loss is capable of producing non-trivial discrepancies in the respective loss landscapes. Above all, it is critical to bear in mind the surrogate loss as another hyperparameter, akin to step-size or optimizer, which has the

$$CE(x, y_1) = -z_1 + \log \sum_j \exp(z_j) \qquad CW(x, y_1) = -z_1 + \max(z_2, z_3)$$
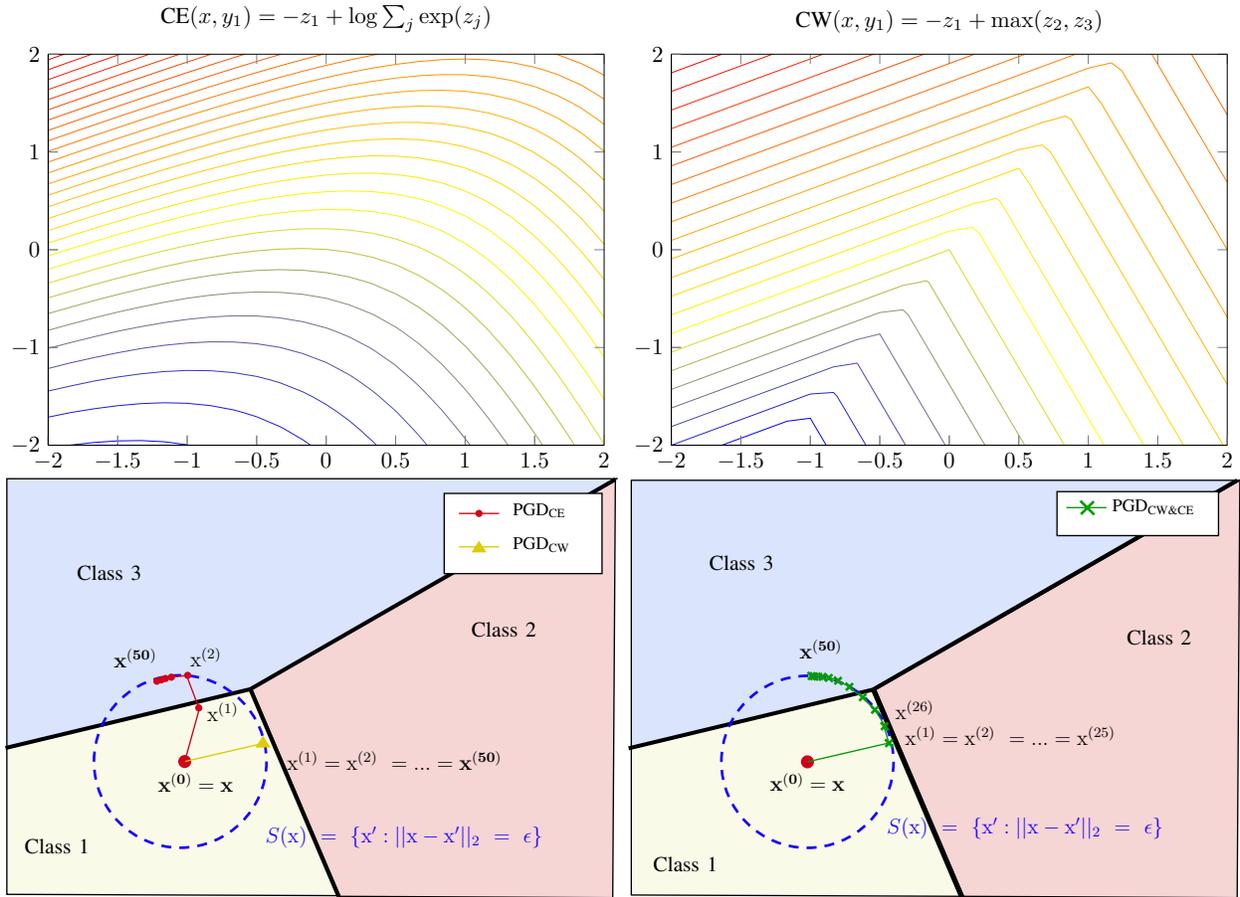
Fig. 1. *Top row*: The level sets of CE and CW losses (w.r.t class $y = 1$). *Bottom row*: (Left) Intermediate PGD points, executed with a single surrogate, where red circles indicate PGD with CE and the yellow triangle PGD with CW, (Right) Intermediate PGD points, but here the objective changes in the middle point ($T = T/2$) of the procedure (green crosses). The blue dashed circle visualizes the boundary surface, which in this case is a disk of radius $\epsilon = 0.4$ centered at $\mathbf{x}$, of the feasible PGD solutions.

potential of causing some degree of robustness overestimation on its own right.

The most straightforward mitigation for this behaviour is to aggregate many different formulations in the same run of PGD. The aggregation of objectives may be instantiated in a variety of ways. Our work is based on a simple *idea* for performing such an aggregation: Divide the PGD process into multiple successive stages, where the surrogate loss changes in the beginning of every stage, and the starting point of every stage coincides with the last step iterate of the previous one. This procedure, when using K stages, can be formulated as:

$$\mathcal{L}(\mathbf{x}, y) = \begin{cases} \mathcal{L}_1(\mathbf{x}, y), & \text{if } t < \frac{T}{K} \\ \mathcal{L}_2(\mathbf{x}, y), & \text{if } \frac{T}{K} \leq t < \frac{2T}{K} \\ \quad \vdots \\ \mathcal{L}_K(\mathbf{x}, y), & \text{if } \frac{(K-1)T}{K} \leq t < T \end{cases}$$

In this paper, we will consider the cases where $K = 2, 3$, using for surrogates the most common choices: CE, CW and DLR.

Notice how this alternation strategy can be viewed as a more complicated initialization: Each PGD stage starts from the initial point $\mathbf{x}^{(0)} = \mathbf{x} + \boldsymbol{\delta}$, where $\boldsymbol{\delta}$: the accumulated perturbation of all previous stages. Of course, an immediate extension is to consider variable starting timesteps t_k for stage k, but in this work, we heuristically set equal time intervals between all stages.

In the remaining discussion, our loss switching variant will be abbreviated as follows: $PGD_{\mathcal{L}_1 \& \mathcal{L}_2 \& \ldots \& \mathcal{L}_K}$, e.g., PGD_{CE} for simple PGD with CE surrogate and $PGD_{CE\&CW}$ for two-stage PGD with CE and CW.

V. EXPERIMENTS

A. Toy Example.

We present a toy example which elucidates that using a single surrogate during PGD may deteriorate performance. Assume a 2D problem of 3-way classification (classes: y_1, y_2, y_3). Inputs are $\mathbf{x} = (x_1, x_2)^T$ and the linear classifier is $\mathbf{z} = (z_1, z_2, z_3)^T = \mathbf{W}\mathbf{x}$, with:

$$\mathbf{W} = \begin{bmatrix} 0.3 & -0.3 \\ 1 & -0.01 \\ -0.25 & 0.75 \end{bmatrix}$$

Dataset	#	Paper	Model ID in RobustBench leaderboard	Architecture	Standard Acc. (%)
CIFAR-10	1	[26]	Engstrom2019Robustness	ResNet-50	87.03
	2	[22]	Carmon2019Unlabeled	WideResNet-28-10	89.69
	3	[36]	Hendrycks2019Using	WideResNet-28-10	87.11
	4	[37]	Zhang2019You	WideResNet-34-10	87.20
	5	[20]	Zhang2019Theoretically	WideResNet-34-10	84.92
	6	[38]	Wu2020Adversarial	WideResNet-34-10	85.36
	7	[39]	Sehwag2021Proxy_R18	ResNet-18	84.59
	8	[19]	Andriushchenko2020Understanding	PreActResNet-18	79.84
	9	[40]	Dai2021Parameterizing	WideResNet-28-10	87.02
	10	[41]	Gowal2021Improving_28_10_ddpm_100m	WideResNet-28-10	87.50
	11	[42]	Huang2021Exploring_ema	WideResNet-34-R	91.23
	12	[43]	Zhang2020Geometry	WideResNet-28-10	89.36
	13	[44]	Rade2021Helper_R18_extra	PreActResNet-18	89.02
	14	[45]	Addepalli2021Towards_RN18	ResNet-18	80.24
	15	[46]	Sehwag2020Hydra	WideResNet-28-10	88.98
CIFAR-100	1	[44]	Rade2021Helper_R18_ddpm	PreActResNet-18	61.50
	2	[47]	Rebuffi2021Fixing_R18_ddpm	PreActResNet-18	56.87
	3	[45]	Addepalli2021Towards_PARN18	PreActResNet-18	62.02
	4	[18]	Rice2020Overfitting	PreActResNet-18	53.83
	5	[36]	Hendrycks2019Using	WideResNet-28-10	59.23
	6	[47]	Rebuffi2021Fixing_28_10_cutmix_ddpm	WideResNet-28-10	62.41
ImageNet	1	[48]	Salman2020Do_R18	ResNet18	52.92
	2	[49]	Salman2020Do_R50	ResNet50	64.02
	3	[26]	Engstrom2019Robustness	ResNet50	62.56
	4	[17]	Wong2020Fast	ResNet50	55.62

TABLE I

OUR MODEL COLLECTION, CONSISTING OF 25 ℓ_∞-BOUNDED DEFENSES OBTAINED FROM THE MODELZOO OF ROBUSTBENCH.

Consider an input $\mathbf{x} = (-0.45, -0.8)$, belonging to the class y_1. The linear model classifies it correctly to its ground-truth class, since $z_1 > \max(z_2, z_3)$. Suppose that our goal is to generate a perturbation δ of bounded ℓ_2-norm (say $\epsilon = 0.4$). A straightforward way to achieve this is by executing PGD, maximizing a surrogate loss, e.g., CE or CW. For the input $\mathbf{x}$ of class y_1, these losses are analytically calculated as:

$$\text{CE}(\mathbf{x}, y) = -z_1 + \log\Big(\sum_{j=1}^{3} \exp(z_j)\Big)$$
$$\text{CW}(\mathbf{x}, y) = -z_1 + \max(z_2, z_3)$$

Figure 1 illustrates the level sets of these two objectives. In the bottom left panel of Figure 1, we visualize the optimization trajectories of PGD for different choices of surrogates. The learning rate is held fixed to $\eta = 2\epsilon$ and PGD is executed for $T = 50$ iterations. The blue dashed circle denotes the boundary of the feasible region, whereas the circle, triangle and cross-shaped points show the intermediate points of PGD $(\mathbf{x}^{(1)}, ..., \mathbf{x}^{(50)})$. Using the CE as surrogate (red circle points) manages to successfully perturb the input $\mathbf{x}$, but CW objective (yellow triangle points) fails because the linear level sets produce gradients that gets the optimization jammed on a single point. The bottom right panel, however, demonstrates that the loss alternation method (green cross points) isn't affected from the failure mode of CW and finds an adversary.

Despite being restricted, this synthetic toy example underpins the argument that using multiple surrogates in the same run of PGD renders the overall procedure more "robust" in the objective selection: Even if some individual choice is infertile

for whatever reason, the other alternatives may be enough to find an adversary.

B. Models and Datasets

We will conduct our experiments in a sizable collection of 25 ℓ_∞-bounded robust models. Specifically, the collection comprises of 15 and 6 defenses on CIFAR-10 and CIFAR-100 [50] respectively, trained with $\epsilon = 8/255$, and 4 defenses on ImageNet [51], trained with perturbation bound $\epsilon = 4/255$. The models are pre-trained and readily obtained from the ModelZoo of RobustBench [52] library. Our collection's robust models originate from various recent works: [18]–[20], [22], [26], [36]–[47], [49]. The architectures of these models are ResNets [53] and Wide ResNets (WRN) [54]. In Table I, we exhibit our model collection: For each case (row), the classifier is matched with the respective paper/work, architecture, ModelID from RobustBench ModelZoo and the accuracy that the classifier attains on the respective clean evaluation set. In the case of CIFAR-10 and CIFAR-100, this coincides with the 10,000 images of the standard test set, whereas in the ImageNet case, 5,000 images from the val set are picked, accordingly to the established split of RobustBench library. We also state that in the following discussion, we'll refer to the terms Attack Success Rate (ASR) and Robust Accuracy (equal to $1-$ASR) interchangeably to quantify the strength of each attack.

C. Experimental Analysis

1) Multi-Stage PGD versus Single-Loss: First, we compare the loss alternation strategy against the typical single loss variants of PGD. In this experimental setting, step size is

		K=1			K=2			K=3
	Model ID	PGD$_{CE}$	PGD$_{CW}$	PGD$_{DLR}$	PGD$_{CE\&CW}$	PGD$_{CE\&DLR}$	PGD$_{CW\&DLR}$	PGD$_{CE\&CW\&DLR}$
CIFAR-10	Engstrom2019Robustness [26]	52.24	52.59	53.55	50.29 -1.95	**50.22** -2.02	52.63 +0.04	50.27 -1.97
	Carmon2019Unlabeled [22]	62.09	60.86	61.16	60.00 -0.86	60.00 -1.16	60.88 +0.02	**59.97** -0.89
	Hendrycks2019Using [36]	57.38	56.61	57.47	55.41 -1.20	55.37 -2.01	56.55 -0.06	**55.35** -1.26
	Zhang2019You [37]	46.28	47.44	47.97	45.33 -0.95	**45.32** -0.96	47.42 -0.02	**45.32** -0.96
	Zhang2019Theoretically [20] †	55.47	54.21	54.39	53.45 -0.76	53.43 -0.96	54.23 +0.02	**53.41** -0.80
	Wu2020Adversarial [38]	59.05	56.93	57.02	56.47 -0.46	56.44 -0.58	56.94 +0.01	**56.42** -0.51
	Sehwag2021Proxy_R18 [39]	58.68	57.22	57.89	56.06 -1.16	**56.05** -1.84	57.21 -0.01	56.06 -1.16
	Andriushchenko2020Understanding [19]	47.14	46.62	47.62	44.56 -2.06	44.53 -2.61	46.62 0	**44.50** -2.12
	Dai2021Parameterizing [40]	63.98	63.23	63.83	61.80 -1.43	**61.76** -2.07	63.23 0	61.77 -1.46
	Gowal2021Improving_28_10_ddpm_100m [41]	65.79	65.20	65.76	63.86 -1.34	63.85 -1.91	65.20 0	**63.84** -1.36
	Huang2021Exploring_ema [42]	64.95	64.15	64.64	63.09 -1.06	**63.03** -1.61	64.12 -0.03	63.06 -1.09
	Zhang2020Geometry [43]	66.67	60.40	60.59	59.78 -0.62	**59.69** -0.90	60.37 -0.03	**59.69** -0.71
	Rade2021Helper_R18_extra [44]	61.48	58.51	58.56	57.77 -0.74	**57.74** -0.82	58.51 0	**57.74** -0.77
	Addepalli2021Towards_RN18 [45]	56.00	51.88	51.97	51.45 -0.43	51.43 -0.54	51.86 -0.03	**51.41** -0.47
	Sehwag2020Hydra [46]	59.86	58.41	58.57	57.66 -0.75	**57.61** -0.96	58.40 -0.01	**57.61** -0.80
CIFAR-100	Rade2021Helper_R18_ddpm [44]	32.60	29.66	29.69	29.12 -0.54	29.08 -0.61	29.66 0	29.12 -0.54
	Rebuffi2021Fixing_R18_ddpm [47]	31.82	29.20	29.25	28.68 -0.52	**28.65** -0.60	29.20 0	28.68 -0.52
	Addepalli2021Towards_PARN18 [45]	32.90	28.10	28.23	27.68 -0.42	**27.63** -0.60	28.10 0	27.67 -0.43
	Rice2020Overfitting [18]	20.89	20.42	20.62	19.33 -1.09	19.33 -1.29	20.42 0	**19.32** -1.10
	Hendrycks2019Using [36]	33.17	30.84	32.34	29.43 -1.41	29.43 -2.91	30.83 -0.01	**29.36** -1.48
	Rebuffi2021Fixing_28_10_cutmix_ddpm [47]	35.74	33.60	33.67	32.53 -1.07	**32.50** -1.17	33.60 0	32.53 -1.07
ImageNet	Salman2020Do_R18 [49]	29.50	27.32	27.60	25.64 -1.68	**25.62** -1.98	27.32 0	25.66 -1.66
	Salman2020Do_R50 [49]	38.78	37.62	38.04	35.30 -2.32	**35.26** -2.78	37.64 +0.02	**35.26** -2.36
	Engstrom2019Robustness [26]	32.64	32.64	33.16	30.00 -2.64	**29.94** -2.70	32.66 +0.02	29.96 -2.68
	Wong2020Fast [17]	27.50	27.46	27.86	25.76 -1.70	**25.74** -1.72	27.48 +0.02	**25.74** -1.72

TABLE II

COMPARING SINGLE-LOSS PGD WITH THE MULTI-STAGE VARIANT OF PGD (WITH $K = 2, 3$). PGD STARTS FROM THE CLEAN POINT (NO ADDED NOISE). THE EXPERIMENTS ARE EXECUTED FOR $T = 100$ WITH NO RESTARTS. EACH ENTRY REPORTS THE ROBUST ACCURACY OF EACH CLASSIFIER FOR THE GIVEN METHOD. (†): ATTACKED WITH $\epsilon = 0.031$. THE GREEN (RED) NUMBERS INDICATE THE RELATIVE DECREASE (INCREASE) OF ROBUST ACCURACY WITH RESPECT TO THE BEST SINGLE-LOSS ATTACK OF EACH MULTI-LOSS VARIANT.

held fixed to $\eta^{(t)} = \epsilon/4$ and the optimizer is set to standard gradient with the sign operation. Our computational budget is $T = 100$ with no restarts. Since no restarts are used, we choose to initiate PGD from the clean points (no initial perturbation) in order to eliminate any source of randomness in the results. Table II presents the robust accuracy obtained of PGD with different choices of surrogates, for every classifier in our collection.

Overall, there are several noteworthy remarks: First, the single-loss columns ($K = 1$) demonstrate that the surrogate loss can greatly affect the ASR of PGD, confirming the findings of previous studies, as that of Croce and Hein [1]. On average, margin loss is the most reliable option but there are cases where it performs worse than CE. There are instances where CE lags behind the other two options by a large margin, e.g., as in the model from [45] (Addepalli2021Towards_RN18), where the gap is greater than 4%. This indicates that it is impossible to select *a priori* the best possible objective for a given model. This observation consitutes strong evidence that no surrogate loss is reliable enough on its own.

Next, the results highlight the advantage of using multiple losses in the same run of PGD: When combining CE with CW or DLR (PGD$_{CE\&CW}$ and PGD$_{CE\&DLR}$ columns), or both (PGD$_{CE\&CW\&DLR}$ column) the attack is always stronger (lower rob. acc.) than the respective single-loss PGD. On average, PGD$_{CE\&CW}$ and PGD$_{CE\&DLR}$ decrease robust accuracy by 1.05% and 1.39% (absolute) respectively, over their corre-

sponding single-loss variants in the CIFAR-10 case. In CIFAR-100, the average absolute decrease in the robust accuracy of the models is 0.81% and 1.19% for PGD$_{CE\&CW}$ and PGD$_{CE\&DLR}$. For the ImageNet dataset, the alternation strategy provides even greater improvements, since the corresponding average reduction reaches 2.08% and 2.29% In the case of PGD$_{CW\&DLR}$, the obtained ASR is nearly identical with PGD$_{CW}$, implying that the alternation step in this case may be futile due to the similarity between the expressions of CW and DLR losses. Overall, our experiments illustrate that the alternation strategy is highly beneficial, across all models and datasets.

Finally, it is illustrated that on average PGD$_{CE\&CW\&DLR}$ is better than PGD$_{CE\&CW}$ and PGD$_{CE\&DLR}$ (mainly on the CIFAR-10 case), yet the differences are small. In some cases, using the alternation scheme with two stages is better than PGD$_{CE\&CW\&DLR}$. This informs us that it is not always better to add another stage/objective in the alternation process. In a fixed iteration budget, adding another loss reduces the overall time allotted to each stage. We assume that this hurts performance because the reduced number of iterations is not enough to reach the stagnating region of each loss.

For the remainder of the experimental section, we will focus on the adversarial defenses of CIFAR-10 dataset.

2) Multi-Stage PGD versus AutoAttack Components: Next, we compare our best method (on average, that is PGD$_{CE\&CW\&DLR}$) with every white-box component from Au-

Model ID	APGD_{CE}	APGD_{DLR}	FAB	$\text{PGD}_{\text{CE\&CW\&DLR}}$	Δ
Engstrom2019Robustness [26]	51.72	52.67	50.67	**50.27**	-0.40
Carmon2019Unlabeled [22]	61.74	60.67	60.88	**59.97**	-0.70
Hendrycks2019Using [36]	57.23	57.03	55.55	**55.35**	-0.20
Zhang2019You [37]	46.15	47.39	45.83	**45.32**	-0.51
Zhang2019Theoretically [20] †	55.28	53.52	53.92	**53.41**	-0.11
Wu2020Adversarial [38]	58.90	56.68	56.82	**56.42**	-0.26
Sehwag2021Proxy_R18 [39]	58.38	57.37	56.27	**56.06**	-0.21
Andriushchenko2020Understanding [19]	46.93	47.08	44.72	**44.50**	-0.22
Dai2021Parameterizing [40]	63.93	63.44	62.27	**61.77**	-0.50
Gowal2021Improving_28_10_ddpm_100m [41]	65.63	65.14	64.14	**63.84**	-0.30
Huang2021Exploring_ema [42]	64.55	64.14	64.45	**63.06**	-1.08
Zhang2020Geometry [43]	66.37	60.19	59.97	**59.69**	-0.28
Rade2021Helper_R18_extra [44]	61.40	58.41	58.42	**57.74**	-0.67
Addepalli2021Towards_RN18 [45]	55.80	51.56	51.93	**51.41**	-0.15
Sehwag2020Hydra [46]	59.60	58.29	58.29	**57.61**	-0.68

TABLE III

COMPARING $\text{PGD}_{\text{CE\&CW\&DLR}}$ WITH THE UNTARGETED VERSION OF EVERY SINGLE WHITE-BOX COMPONENT FROM THE AUTOATTACK ENSEMBLE. EACH ENTRY REPORTS THE ROBUST ACCURACY OF EACH CLASSIFIER FOR THE GIVEN METHOD. Δ COLUMN REPORT THE ROBUST ACCURACY GAP BETWEEN $\text{PGD}_{\text{CE\&CW\&DLR}}$ AND THE BEST AMONG THE AUTOATTACK COMPONENTS. THE EXPERIMENTS ARE EXECUTED FOR $T = 100$ WITH NO RESTARTS. (†): ATTACKED WITH $\epsilon = 0.031$.

Model ID	GAMA-PGD [13]	$\text{PGD}_{\text{CE\&CW\&DLR}}$ (GAMA-PGD sch.)	Δ	MD Attack [11]	$\text{PGD}_{\text{CE\&CW\&DLR}}$ (MD sch.)	Δ
Engstrom2019Robustness [26]	50.05	**49.88**	-0.17	50.34	**49.87**	-0.47
Carmon2019Unlabeled [22]	59.84	**59.78**	-0.06	59.83	**59.72**	-0.11
Hendrycks2019Using [36]	**55.22**	55.26	+0.04	**55.15**	55.20	+0.05
Zhang2019You [37]	45.32	**45.20**	-0.12	45.49	**45.17**	-0.32
Zhang2019Theoretically [20]†	53.29	53.29	0	53.36	**53.26**	-0.10
Wu2020Adversarial [38]	56.30	56.30	0	56.28	**56.26**	-0.02
Sehwag2021Proxy_R18 [39]	56.01	**55.95**	-0.06	55.92	**55.89**	-0.03
Andriushchenko2020Understanding [19]	44.42	**44.41**	-0.01	44.57	**44.44**	-0.13
Dai2021Parameterizing [40]	61.94	**61.74**	-0.20	61.99	**61.72**	-0.27
Gowal2021Improving_28_10_ddpm_100m [41]	63.78	**63.72**	-0.06	63.94	**63.73**	-0.21
Huang2021Exploring_ema [42]	**62.87**	62.89	+0.02	62.93	**62.86**	-0.07
Zhang2020Geometry [43]	60.72	**59.62**	-1.10	59.73	**59.58**	-0.15
Rade2021Helper_R18_extra [44]	57.78	**57.73**	-0.05	57.74	**57.72**	-0.02
Addepalli2021Towards_RN18 [45]	51.43	**51.26**	-0.17	51.30	**51.25**	-0.05
Sehwag2020Hydra [46]	57.49	**57.43**	-0.06	**57.31**	57.45	+0.14

TABLE IV

COMPARING $\text{PGD}_{\text{CE\&CW\&DLR}}$ WITH THE STRONGEST ATTACKS OF OUR COMPUTATIONAL BUDGET ($\mathbf{T = 100}, \mathbf{R = 1}$). EACH ENTRY REPORTS THE ROBUST ACCURACY OF EACH CLASSIFIER FOR THE GIVEN METHOD. Δ COLUMNS REPORT THE ROBUST ACCURACY GAP BETWEEN THE COMPARED METHODS. (†): ATTACKED WITH $\epsilon = 0.031$.

Model ID	GAMA-PGD [13]	$\text{PGD}_{\text{CE\&CW\&DLR}}$ (GAMA-PGD sch.)	Δ	MD Attack [11]	$\text{PGD}_{\text{CE\&CW\&DLR}}$ (MD sch.)	Δ
Engstrom2019Robustness [26]	49.88	**49.80**	-0.08	50.13	**49.68**	-0.45
Carmon2019Unlabeled [22]	**59.70**	59.71	+0.01	59.67	59.67	0
Hendrycks2019Using [36]	55.21	**55.09**	-0.12	**55.10**	55.13	+0.03
Zhang2019You [37]	45.02	45.02	0	45.36	**45.00**	-0.36
Zhang2019Theoretically [20]†	53.26	**53.20**	-0.06	53.26	**53.19**	-0.07
Wu2020Adversarial [38]	56.29	**56.23**	-0.06	56.24	**56.20**	-0.04
Sehwag2021Proxy_R18 [39]	55.89	**55.85**	-0.04	55.85	**55.81**	-0.04
Andriushchenko2020Understanding [19]	44.32	**44.29**	-0.03	44.42	**44.36**	-0.06
Dai2021Parameterizing [40]	61.95	**61.70**	-0.25	61.83	**61.65**	-0.18
Gowal2021Improving_28_10_ddpm_100m [41]	63.81	**63.66**	-0.15	63.90	**63.65**	-0.25
Huang2021Exploring_ema [42]	62.77	**62.74**	-0.03	62.80	**62.71**	-0.09
Zhang2020Geometry [43]	60.28	**59.44**	-0.88	59.59	**59.50**	-0.09
Rade2021Helper_R18_extra [44]	57.74	**57.66**	-0.08	57.71	**57.67**	-0.04
Addepalli2021Towards_RN18 [45]	51.46	**51.23**	-0.23	51.22	**51.21**	-0.01
Sehwag2020Hydra [46]	57.37	**57.31**	-0.06	**57.23**	57.34	+0.11

TABLE V

COMPARING $\text{PGD}_{\text{CE\&CW\&DLR}}$ WITH THE SIMILAR (TO OUR WORK) BASELINES ($\mathbf{T = 100}, \mathbf{R = 5}$). EACH ENTRY REPORTS THE ROBUST ACCURACY OF EACH CLASSIFIER FOR THE GIVEN METHOD. Δ COLUMNS REPORT THE ROBUST ACCURACY GAP BETWEEN THE COMPARED METHODS. (†): ATTACKED WITH $\epsilon = 0.031$.
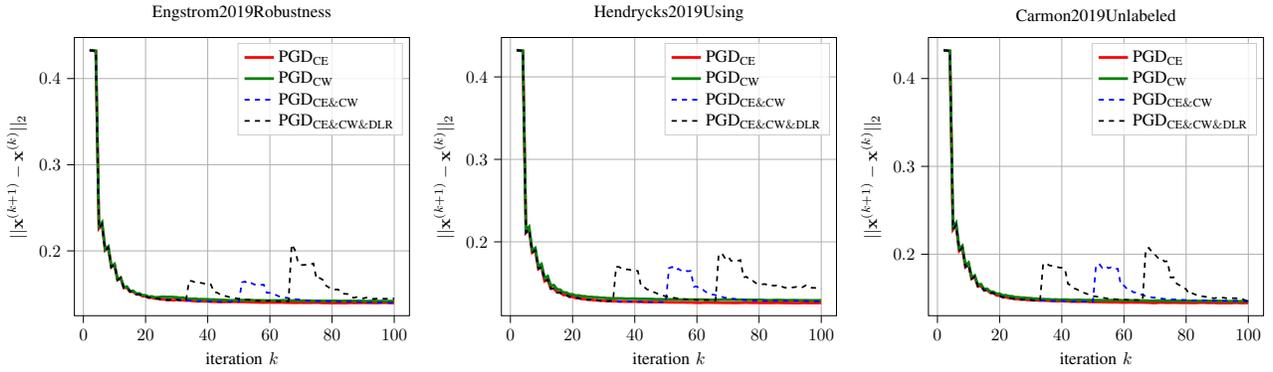
Fig. 2. Plotting the ℓ_2-norm between successive PGD steps, for various surrogate losses. Each panel represents this quantity over iterations, for a different classifier (ModelID is on top of each panel).

toAttack [1], i.e., APGD$_{CE}$, APGD$_{DLR}$ and FAB attack [12]. In the original AutoAttack evaluation, the last two components are run for $T = 100$ iterations and $R = 9$ restarts, using the targeted version of each attack. However, we adapt these attacks to our computational budget, evaluating the performance of their untargeted versions for $T = 100$. In our experiments, we execute the official code[1] of AutoAttack for every single model. We clarify that the official code does not provide a way to turn off random initialization when evaluating the AA components, but the fluctuations are expected to be small enough.

As it is clearly illustrated in Table III, our proposed method, PGD$_{CE\&CW\&DLR}$ consistently outperforms the white-box components of AutoAttack. It becomes evident that the advantage of using the loss switching strategy is significant, since in this setting we run our attack for fixed step size equal to $\epsilon/4$ and the simplest optimizer possible (sign operation with no momentum). APGD$_{CE}$ and APGD$_{DLR}$ are both based in the evidently better APGD optimizer and step size is decayed according to some schedule, yet they lag behind PGD$_{CE\&CW\&DLR}$ by a large margin. Particularly, PGD$_{CE\&CW\&DLR}$ achieves (on average) 0.418% lower robust accuracy than the strongest component.

3) Multi-Stage PGD versus the strongest baselines: We extend the assessment of our method's effectiveness by comparing it with the strongest ℓ_∞-bounded attacks on CIFAR-10, for $T = 100$ and no restarts. We consider the two best baselines found in literature (in our computational budget): GAMA-PGD [13] and MD attack [11]. Both of these methods suggest improving PGD through modifications on the surrogate loss and step size schedule. In Subsection VI-A, we delve into the exact similarities between the examined methods and our work.

We execute these attacks through the official codebases[2] [3]. When comparing PGD$_{CE\&CW\&DLR}$ with each baseline, we adapt the learning rate schedule according to each work

[1]https://github.com/fra31/auto-attack
[2]https://github.com/val-iisc/GAMA-GAT
[3]https://github.com/Jack-lx-jiang/MD_attacks

(See Appendix for details). The results of these comparisons are summarized in Table IV. In the parentheses of PGD$_{CE\&CW\&DLR}$ columns, we display which learning rate schedule is used. These results indicate the effectiveness of our attack, achieving state-of-the-art performance (in the $T = 100$, $R = 1$ budget), for the majority of evaluated models.

Specifically, PGD$_{CE\&CW\&DLR}$ outperforms GAMA-PGD [13] in 11 out of 15 ℓ_∞-bounded robust models, whereas in 2 models they achieve the exact same ASR. In the 2 networks that PGD$_{CE\&CW\&DLR}$ returns higher robust accuracy, the differences are quite small, i.e., 0.02% and 0.04%. An extreme case is the model of [43], since GAMA-PGD lags behind our method for 1.10%. These observations indicate that, in general, PGD$_{CE\&CW\&DLR}$ suffers less from robustness overestimation.

In the case of MD attack [11], our method achieves lower robust accuracy in 13 out of 15 tested models, with an average improvement of 0.15%. In two models [36], [46], however, the estimated robust accuracy is 0.05% and 0.14% higher than that of MD attack. Overall, this comparison, similarly to the previous one, highlights that PGD$_{CE\&CW\&DLR}$ provides the most reliable ℓ_∞-bounded robustness evaluations.

Since the differences of our best method with these baselines are marginal for some cases, it is crucial to answer whether they arise just because our method is benefited from the specific PGD starting point (which in our case is the clean image). To address this, we repeat the above comparisons for the same amount of iterations but with $R = 5$ restarts. In this case, the starting points are initialized with random noise $\delta = \epsilon \cdot \text{sgn}(\mathbf{u})$, where $\mathbf{u} \sim \mathcal{U}(-1, 1)$. The results are illustrated in Table V. Overall, it is evident that the increased number of restarts helps each attack to achieve slightly lower robust accuracy, but comparatively, our attack still performs more reliable robustness evaluations for the vast majority of cases.

D. Qualitative Analysis

Here, we conduct a qualitative analysis to better grasp the impact of changing surrogate losses during optimization. Our experiments are inspired by the work of Yamamura et al. [33], where they visualize the ℓ_2-distance between successive

Model ID	PGD$_{CE\&CW}$	PGD$_{CW\&CE}$	PGD$_{CE\&DLR}$	PGD$_{DLR\&CE}$	PGD$_{CW\&DLR}$	PGD$_{DLR\&CW}$
Engstrom2019Robustness [26]	**50.29**	50.95	**50.22**	51.13	**52.63**	52.97
Carmon2019Unlabeled [22]	**60.00**	60.27	**60.00**	60.39	**60.88**	60.94
Hendrycks2019Using [36]	**55.41**	55.62	**55.37**	55.72	**56.55**	56.84
Zhang2019You [37]	**45.33**	45.85	**45.32**	45.86	**47.42**	47.58
Zhang2019Theoretically [20]	**53.45**	53.76	**53.43**	53.86	**54.23**	54.31
Wu2020Adversarial [38]	**56.47**	56.68	**56.44**	56.72	**56.94**	56.98
Sehwag2021Proxy_R18 [39]	**56.06**	56.42	**56.05**	56.52	**57.21**	57.49
Andriushchenko2020Understanding [19]	**44.56**	44.94	**44.53**	45.03	**46.62**	46.77
Dai2021Parameterizing [40]	**61.80**	62.18	**61.76**	62.33	**63.23**	63.42
Gowal2021Improving_28_10_ddpm_100m [41]	**63.86**	64.80	**63.85**	64.30	**65.20**	65.31
Huang2021Exploring_ema [42]	**63.09**	63.52	**63.03**	63.64	**64.12**	64.34
Zhang2020Geometry [43]	**59.78**	60.16	**59.69**	60.31	**60.37**	60.51
Rade2021Helper_R18_extra [44]	**57.77**	58.17	**57.74**	58.18	**58.51**	58.54
Addepalli2021Towards_RN18 [45]	**51.45**	51.79	**51.43**	51.84	**51.86**	51.93
Sehwag2020Hydra [46]	**57.66**	57.92	**57.61**	57.93	**58.40**	58.47

TABLE VI

ABLATION STUDY. EXPLORING THE IMPORTANCE OF THE SURROGATES' ORDER. THE EXPERIMENTS ARE EXECUTED FOR $T = 100$ WITH NO RESTARTS. EACH ENTRY REPORTS THE ROBUST ACCURACY OF EACH CLASSIFIER FOR THE GIVEN METHOD. (†): ATTACKED WITH $\epsilon = 0.031$.

				Convex		
Model ID	PGD$_{CE}^{100}$	PGD$_{CW}^{100}$	PGD$_{CE}^{50} \vee$ PGD$_{CW}^{50}$	$\gamma = 0.25$	$\gamma = 0.75$	PGD$_{CE\&CW}$
Engstrom2019Robustness [26]	52.24	52.59	50.75	51.59	52.30	**50.29**
Carmon2019Unlabeled [22]	62.09	60.86	60.18	60.97	60.85	**60.00**
Hendrycks2019Using [36]	57.38	56.61	55.52	56.10	56.36	**55.41**
Zhang2019You [37]	46.28	47.44	45.56	46.37	47.19	**45.33**
Zhang2019Theoretically [20] †	55.47	54.21	53.68	54.34	54.18	**53.45**
Wu2020Adversarial [38]	59.05	56.93	56.66	57.46	57.00	**56.47**
Sehwag2021Proxy_R18 [39]	58.68	57.22	56.37	57.22	57.10	**56.06**
Andriushchenko2020Understanding [19]	47.14	46.62	44.81	45.78	46.17	**44.56**
Dai2021Parameterizing [40]	63.98	63.23	62.17	62.89	63.15	**61.80**
Gowal2021Improving_28_10_ddpm_100m [41]	65.79	65.20	64.20	64.72	65.00	**63.86**
Huang2021Exploring_ema [42]	64.95	64.15	63.35	64.01	64.09	**63.09**
Zhang2020Geometry [43]	66.67	60.40	60.14	63.88	60.86	**59.78**
Rade2021Helper_R18_extra [44]	61.48	58.51	58.14	59.14	58.49	**57.77**
Addepalli2021Towards_RN18 [45]	56.00	51.88	51.78	53.23	51.96	**51.45**
Sehwag2020Hydra [46]	59.86	58.41	57.85	58.59	58.41	**57.66**

TABLE VII

ABLATION STUDY. IN THE CONVEX COLUMNS, γ $(1 - \gamma)$ CORRESPONDS TO CE (CW). THE EXPERIMENTS ARE EXECUTED FOR $T = 100$ WITH NO RESTARTS. EACH ENTRY REPORTS THE ROBUST ACCURACY OF EACH CLASSIFIER FOR THE GIVEN METHOD. (†): ATTACKED WITH $\epsilon = 0.031$.

PGD steps: $\|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\|_2$ in order to empirically show that their proposed optimizer explores the input space more extensively. In a similar vein, we replicate their method for PGD$_{CE}$, PGD$_{CW}$, PGD$_{CE\&CW}$, PGD$_{CE\&CW\&DLR}$ in Figure 2, inspecting three different classifiers. To generate smoother curves, the y-axis quantity is averaged on a batch of 100 examples.

Altogether, it appears that in the single loss variants, the search of PGD becomes quite localized and after some time the successive steps are within small distances. In the cases where multiple surrogates are used, the curve presents a sudden rise in the alternation timestep, indicating that the objective alternation helps the algorithm to diversify its search.

E. Ablation: Surrogate Loss Order in Multi-Stage PGD

A research question regarding the multi-stage variant of PGD is whether the objective ordering affects the results. Specifically, we are interested in understanding whether any change occurs if we optimize the objectives with reverse ordering. To address this question, we execute the two-stage

PGD, with $T = 100$ and no restarts, for every possible pair (order matters) of CE, CW and DLR.

The results of Table VI demonstrate that the order plays an essential role. Particularly, it is clearly illustrated that it is better to start the optimization procedure with the CE loss, then finishing off with CW or DLR. However, we observe that regardless of the objective ordering, every multi-stage PGD variant which alternates between CE and one of CW, DLR (PGD$_{CE\&CW}$, PGD$_{CW\&CE}$, PGD$_{CE\&DLR}$, PGD$_{DLR\&CE}$ columns) performs better than single-loss PGD.

F. Ablation: Additional Techniques of Combining Surrogates

Another interesting research question is to explore whether there exist other ways of combining surrogates. To settle this, we compare the alternation method with two additional combining techniques. First, one can combine different surrogates through a convex combination, i.e., setting the surrogate according to the following expression:

$$\mathcal{L}(\mathbf{x}, y) = \gamma \cdot \mathcal{L}_1(\mathbf{x}, y) + (1 - \gamma) \cdot \mathcal{L}_2(\mathbf{x}, y)$$

Another way is to combine different surrogates in an ensemble-like manner, i.e., split the entire iteration budget into K equally sized intervals, execute PGD using the k-th surrogate $\mathcal{L}_k$, starting from the clean point (not from where the previous stage ended), and then aggregate the output decisions. This method is inspired by the MultiTargeted surrogate, introduced by Gowal et al. [30]. For the CE and CW losses, we denote the latter combining strategy as $PGD_{CE} \vee PGD_{CW}$, because the output decisions of each surrogate are aggregated through binary OR, i.e., the input is deemed misclassified if at least one of PGD_{CE}, PGD_{CW} generate a successful perturbation.

We conduct an ablation study, using the CE and CW objectives, to explore the effectiveness of these methods. The results are illustrated in Table VII, where we also report the robust accuracy of $PGD_{CE}, PGD_{CW}, PGD_{CE\&CW}$ for direct comparison (we also include the iteration budget on the superscript to draw a distinction with the ensemble method). As expected, the robust accuracy of convex combination is susceptible to the choice of γ, with its performance depending on whether the best-performing objective has a larger weight. The ensemble method, on the other hand, consistently outperforms the single-loss PGD, and much like $PGD_{CE\&CW}$, is more "robust" against issues arising from individual use of objectives. However, the loss alternation strategy, $PGD_{CE\&CW}$, performs better than the ensemble-like combination. We advocate that this occurs because $PGD_{CE\&CW}$ utilizes the progress made in previous stages to perform better initialization for the next stage. The ensemble-like method, however, discards the perturbation found by previous objectives, and starts optimization all over again.

VI. DISCUSSION

A. Similarity with Previous Works

Next, we discuss previous works that also employ a loss alternating strategy. First, the most similar work is that of Ma et al. [11], where they employ an identical alternation step to evade the issue of imbalanced gradients. The first PGD stage optimize only one of the two logit terms, whereas in the final stage, the typical margin loss is optimized. Notice a striking difference: Our work involves the CE,CW and DLR losses, all containing more than one logit terms, hence potentially suffering from gradient imbalance that should translate to reduced ASR. Our method outperforms MD attack. Therefore, our study implies that the performance improvement of MD attack [11] may be the outcome of switching surrogates, rather than deterring the magnitudes of logit terms' gradients from becoming highly disparate.

The second method is GAMA-PGD, introduced by Sriramanan et al. [13]. The authors propose to regularize the margin loss with a MSE term, weighted by a decaying coefficient. In their implementation, the initial rate of weights between the MSE and CW losses is 50:1, hence for the first few iterations the contribution of CW loss is negligible. The weight of MSE is linearly decayed to 0 for $T/4$ iterations, and after that point the surrogate is set to the standard margin loss. Essentially,

their attack alternates the surrogate loss used by PGD as many times as the duration of the interval during which MSE decays, i.e., $T/4$ out of T iterations. Their analysis conveys the intuition that the improvement originates solely from the regularizing effect that MSE exerts on the margin loss. Our work demonstrates that the benefits of GAMA-PGD may arise from the loss alternation, still further experimentation is required.

Another method loosely connected with ours is the Composite Adversarial Attack (CAA) [55]. Mao et al. propose to generate adversaries by searching for the best composition of individual base attacks. Our method can be seen as a more special study of CAA, since it composes PGD attacks for two (or three) different objectives. Our work indicates much more markedly the value of using multiple losses. The effectiveness of CAA appears more like the result of a brute-force-like search.

Overall, our paper differs from the aforementioned works in that it manages to showcase the true efficacy of the alternation step, stripped down from other redundant components. The experiments provide direct evidence that using multiple objectives is sufficient to induce large performance gains. Additionally, our work is an extension of these methods since we evaluate the combination of all possible pairs of CE,CW and DLR losses, rather than using only CW with its individual terms [11] or CW and MSE [13].

B. Future Work

There are several questions arising from the proposed work than require further investigation and could be of value to the community. Notably, it is critical to address whether there is a trade-off between the number of surrogates used and PGD performance, for a fixed number of iterations. We assumed that adding more stages for fixed budget may hinder performance due to the decreased duration allotted to each stage. However, our intuition is that adding more objectives shouldn't drop the Attack Success Rate (ASR), given that PGD spends a sufficient time in each stage. This can be easily verified by increasing the computational budget and observing wether the larger amount of surrogates leads to higher ASR.

Another interesting observation to explore is how the alternation step depends on the choice of objectives and their respective formulations. Particularly, we observed that $PGD_{CW\&DLR}$ performs at a par (or even worse) than the respective single-loss variants, PGD_{CW} and PGD_{DLR}, which was credited to the similarity of CW and DLR. This indicates that the loss alternation technique is an improvement only if the expressions generate landscapes which are diverse enough. In this vein, it would be valuable to encompass other expressions which deviate from the objective functions of our study, i.e., CE, CW and DLR.

Since we experimentally demonstrate that our PGD variant is the strongest adversarial attack in the computational budget of 100 iterations, another direct extension is to integrate our attack into powerful ensembles. Specifically, in the case of AutoAttack [1], $PGD_{CE\&CW\&DLR}$ is outperforming every white-

box component (Table III), hence we assume that replacing e.g. APGD$_{DLR}$ with PGD$_{CE\&CW\&DLR}$ would produce more reliable robustness evaluations.

Apart from that, it is worthwhile to investigate whether the idea of increasing the number of surrogates helps other algorithms to perform better. Notice that our work is entirely framed within the PGD algorithm, but other popular attacks remain unexplored. Subsequent works could address whether our findings extrapolate to other attacks, and even in other settings e.g. black-box attacks.

VII. CONCLUSION

In this work, we propose a method of alternating objectives for improving the strength of PGD-based attacks. The proposed method performs better than single loss variants, across 25 adversarial defenses, spanning 3 different datasets. In the CIFAR-10 case, it performs better than strong baselines which are used for evaluating the ℓ_p-bounded robustness of neural networks: AutoPGD [1], FAB [12], GAMA-PGD [13] and MD Attack [11]. Our experiments show that alternating objectives is a very effective way of combining different objectives compared, e.g., to convex combination and ensemble-like methods. It is also experimentally shown that the proposed method offers significant robustness towards overcoming loss-specific weaknesses. Furthermore, our qualitative analysis offers intuition on the reasons behind our method's strength that may be related to the algorithm's search space diversification induced by the alternation step. Finally, we offer a new perspective on how the success of other state-of-the-art attacks, i.e., GAMA-PGD and MD Attack, can be ascribed to loss alternation.

REFERENCES

[1] F. Croce and M. Hein, "Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks," in *International conference on machine learning*. PMLR, 2020, pp. 2206–2216.

[2] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.

[3] A. Athalye, N. Carlini, and D. A. Wagner, "Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples," in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 2018, pp. 274–283.

[4] J. Uesato, B. O'donoghue, P. Kohli, and A. Oord, "Adversarial risk and the dangers of evaluating against weak attacks," in *International Conference on Machine Learning*. PMLR, 2018, pp. 5025–5034.

[5] F. Tramèr, N. Carlini, W. Brendel, and A. Madry, "On adaptive attacks to adversarial example defenses," in *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

[6] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

[7] A. Kurakin, I. J. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*, 2017.

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

[9] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.

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

[11] X. Ma, L. Jiang, H. Huang, Z. Weng, J. Bailey, and Y.-G. Jiang, "Imbalanced gradients: A subtle cause of overestimated adversarial robustness," *arXiv preprint arXiv:2006.13726*, 2020.

[12] F. Croce and M. Hein, "Minimally distorted adversarial examples with a fast adaptive boundary attack," in *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, ser. Proceedings of Machine Learning Research, vol. 119. PMLR, 2020, pp. 2196–2205.

[13] G. Sriramanan, S. Addepalli, A. Baburaj *et al.*, "Guided adversarial attack for evaluating and enhancing adversarial defenses," *Advances in Neural Information Processing Systems*, vol. 33, pp. 20 297–20 308, 2020.

[14] J. Rony, L. G. Hafemann, L. Oliveira, I. B. Ayed, R. Sabourin, and E. Granger, "Decoupling direction and norm for efficient gradient-based l2 adversarial attacks and defenses," *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4317–4325, 2019.

[15] M. Pintor, F. Roli, W. Brendel, and B. Biggio, "Fast minimum-norm adversarial attacks through adaptive norm constraints," *Advances in Neural Information Processing Systems*, vol. 34, pp. 20 052–20 062, 2021.

[16] A. Shafahi, M. Najibi, A. Ghiasi, Z. Xu, J. P. Dickerson, C. Studer, L. S. Davis, G. Taylor, and T. Goldstein, "Adversarial training for free!" in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 2019, pp. 3353–3364.

[17] E. Wong, L. Rice, and J. Z. Kolter, "Fast is better than free: Revisiting adversarial training," in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, 2020.

[18] L. Rice, E. Wong, and J. Z. Kolter, "Overfitting in adversarially robust deep learning," in *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, ser. Proceedings of Machine Learning Research, vol. 119. PMLR, 2020, pp. 8093–8104.

[19] M. Andriushchenko and N. Flammarion, "Understanding and improving fast adversarial training," *Advances in Neural Information Processing Systems*, vol. 33, pp. 16 048–16 059, 2020.

[20] H. Zhang, Y. Yu, J. Jiao, E. Xing, L. El Ghaoui, and M. Jordan, "Theoretically principled trade-off between robustness and accuracy," in *International conference on machine learning*. PMLR, 2019, pp. 7472–7482.

[21] L. Schmidt, S. Santurkar, D. Tsipras, K. Talwar, and A. Madry, "Adversarially robust generalization requires more data," in *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 2018, pp. 5019–5031.

[22] Y. Carmon, A. Raghunathan, L. Schmidt, J. C. Duchi, and P. S. Liang, "Unlabeled data improves adversarial robustness," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[23] R. Zhai, T. Cai, D. He, C. Dan, K. He, J. E. Hopcroft, and L. Wang, "Adversarially robust generalization just requires more unlabeled data," *ArXiv*, vol. abs/1906.00555, 2019.

[24] J. Alayrac, J. Uesato, P. Huang, A. Fawzi, R. Stanforth, and P. Kohli, "Are labels required for improving adversarial robustness?" in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 2019, pp. 12 192–12 202.

[25] F. Tramèr and D. Boneh, "Adversarial training and robustness for multiple perturbations," in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing*

Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada, 2019, pp. 5858–5868.

[26] L. Engstrom, A. Ilyas, H. Salman, S. Santurkar, and D. Tsipras, "Robustness (python library)," 2019. [Online]. Available: https://github.com/MadryLab/robustness

[27] M. Andriushchenko, F. Croce, N. Flammarion, and M. Hein, "Square attack: a query-efficient black-box adversarial attack via random search," in *European Conference on Computer Vision*. Springer, 2020, pp. 484–501.

[28] Y. Liu, Y. Cheng, L. Gao, X. Liu, Q. Zhang, and J. Song, "Practical evaluation of adversarial robustness via adaptive auto attack," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 15 105–15 114.

[29] Y. Yu, X. Gao, and C.-Z. Xu, "Lafeat: piercing through adversarial defenses with latent features," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 5735–5745.

[30] S. Gowal, J. Uesato, C. Qin, P.-S. Huang, T. A. Mann, and P. Kohli, "An alternative surrogate loss for pgd-based adversarial testing," *ArXiv*, vol. abs/1910.09338, 2019.

[31] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

[32] B. Polyak, "Some methods of speeding up the convergence of iteration methods," *USSR Computational Mathematics and Mathematical Physics*, vol. 4, no. 5, pp. 1–17, 1964.

[33] K. Yamamura, H. Sato, N. Tateiwa, N. Hata, T. Mitsutake, I. Oe, H. Ishikura, and K. Fujisawa, "Diversified adversarial attacks based on conjugate gradient method," in *International Conference on Machine Learning*. PMLR, 2022, pp. 24 872–24 894.

[34] Y. Tashiro, Y. Song, and S. Ermon, "Diversity can be transferred: Output diversification for white-and black-box attacks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 4536–4548, 2020.

[35] S. Arora, L. Babai, J. Stern, and Z. Sweedyk, "The hardness of approximate optima in lattices, codes, and systems of linear equations," *Journal of Computer and System Sciences*, vol. 54, no. 2, pp. 317–331, 1997.

[36] D. Hendrycks, K. Lee, and M. Mazeika, "Using pre-training can improve model robustness and uncertainty," in *International Conference on Machine Learning*. PMLR, 2019, pp. 2712–2721.

[37] D. Zhang, T. Zhang, Y. Lu, Z. Zhu, and B. Dong, "You only propagate once: Accelerating adversarial training via maximal principle," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[38] D. Wu, S.-T. Xia, and Y. Wang, "Adversarial weight perturbation helps robust generalization," *Advances in Neural Information Processing Systems*, vol. 33, pp. 2958–2969, 2020.

[39] V. Sehwag, S. Mahloujifar, T. Handina, S. Dai, C. Xiang, M. Chiang, and P. Mittal, "Robust learning meets generative models: Can proxy distributions improve adversarial robustness?" in *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*, 2022.

[40] S. Dai, S. Mahloujifar, and P. Mittal, "Parameterizing activation functions for adversarial robustness," in *43rd IEEE Security and Privacy, SP Workshops 2022, San Francisco, CA, USA, May 22-26, 2022*. IEEE, 2022.

[41] S. Gowal, S.-A. Rebuffi, O. Wiles, F. Stimberg, D. A. Calian, and T. A. Mann, "Improving robustness using generated data," *Advances in Neural Information Processing Systems*, vol. 34, pp. 4218–4233, 2021.

[42] H. Huang, Y. Wang, S. Erfani, Q. Gu, J. Bailey, and X. Ma, "Exploring architectural ingredients of adversarially robust deep neural networks," *Advances in Neural Information Processing Systems*, vol. 34, pp. 5545–5559, 2021.

[43] J. Zhang, J. Zhu, G. Niu, B. Han, M. Sugiyama, and M. S. Kankanhalli, "Geometry-aware instance-reweighted adversarial training," in *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021.

[44] R. Rade and S.-M. Moosavi-Dezfooli, "Helper-based adversarial training: Reducing excessive margin to achieve a better accuracy vs. robustness trade-off," in *ICML 2021 Workshop on Adversarial Machine Learning*, 2021.

[45] S. Addepalli, S. Jain, G. Sriramanan, S. Khare, and V. B. Radhakrishnan, "Towards achieving adversarial robustness beyond perceptual limits," in *ICML 2021 Workshop on Adversarial Machine Learning*, 2021.

[46] V. Sehwag, S. Wang, P. Mittal, and S. Jana, "Hydra: Pruning adversarially robust neural networks," *Advances in Neural Information Processing Systems*, vol. 33, pp. 19 655–19 666, 2020.

[47] S.-A. Rebuffi, S. Gowal, D. A. Calian, F. Stimberg, O. Wiles, and T. Mann, "Fixing data augmentation to improve adversarial robustness," *arXiv preprint arXiv:2103.01946*, 2021.

[48] H. Salman, G. Yang, J. Li, P. Zhang, H. Zhang, I. P. Razenshteyn, and S. Bubeck, "Provably robust deep learning via adversarially trained smoothed classifiers," in *NeurIPS*, 2019.

[49] H. Salman, A. Ilyas, L. Engstrom, A. Kapoor, and A. Madry, "Do adversarially robust imagenet models transfer better?" *ArXiv*, vol. abs/2007.08489, 2020.

[50] A. Krizhevsky, "Learning multiple layers of features from tiny images," 2009.

[51] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*. IEEE Computer Society, 2009, pp. 248–255.

[52] F. Croce, M. Andriushchenko, V. Sehwag, E. Debenedetti, N. Flammarion, M. Chiang, P. Mittal, and M. Hein, "Robustbench: a standardized adversarial robustness benchmark," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2021.

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

[54] S. Zagoruyko and N. Komodakis, "Wide residual networks," in *Proceedings of the British Machine Vision Conference 2016, BMVC 2016, York, UK, September 19-22, 2016*. BMVA Press, 2016.

[55] X. Mao, Y. Chen, S. Wang, H. Su, Y. He, and H. Xue, "Composite adversarial attacks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 10, 2021, pp. 8884–8892.

APPENDIX

A. Implementation Details

For our experiments, we implement code on the Py-Torch framework. The PGD implementation is based on the TRADES [20] repository[4]. All attacks are executed with a ℓ_∞-norm bound of $\epsilon = 8/255$ and for $T = 100$ iterations, with no restarts. Our code returns the best intermediate PGD point instead of the last. The robust models of our study are obtained from the ModelZoo of RobustBench [52]. Our experiments are run in a NVIDIA GeForce GTX 1080 Ti GPU with 12GB VRAM.

B. Step Size Schedules

Here, we discuss the step size schedules used when comparing our method with the GAMA-PGD [13] and MD Attack [11] baselines. In GAMA-PGD, the step size schedule incurs tenfold drops at $T = 60$ and $T = 85$, starting from $\eta^{(0)} = 2\epsilon$.

In [11], step size is regulated according to a cosine-annealing scheme. In particular, the step size in $t-th$ iteration equals:

$$\eta^{(t)} = \begin{cases} \epsilon \cdot (1 + \cos(\frac{t-1}{T'})\pi) & , t < T' \\ \epsilon \cdot (1 + \cos(\frac{t-T'}{T-T'}\pi)) & , T' \leq t < T \end{cases}$$

where $T = 100$, $T' = T/2$. Therefore, step size is decayed from 2ϵ to 0 in each stage. We extend this scheme to our three-stage variant as follows:

$$\eta^{(t)} = \begin{cases} \epsilon \cdot (1 + \cos(\frac{t-1}{T/3})\pi) & , t < T/3 \\ \epsilon \cdot (1 + \cos(\frac{t-T/3}{T/3}\pi)) & , T/3 \leq t < 2T/3 \\ \epsilon \cdot (1 + \cos(\frac{t-2T/3}{T/3}\pi)) & , 2T/3 \leq t < T \end{cases}$$

[4]https://github.com/yaodongyu/TRADES