

Reliably fast adversarial training via latent adversarial perturbation

Geon Yeong Park Sang Wan Lee

KAIST

Daejeon, South Korea

{pky3436, sangwan}@kaist.ac.kr

Abstract

While multi-step adversarial training is widely popular as an effective defense method against strong adversarial attacks, its computational cost is notoriously expensive, compared to standard training. Several single-step adversarial training methods have been proposed to mitigate the above-mentioned overhead cost; however, their performance is not sufficiently reliable depending on the optimization setting. To overcome such limitations, we deviate from the existing input-space-based adversarial training regime and propose a single-step latent adversarial training method (SLAT), which leverages the gradients of latent representation as the latent adversarial perturbation. We demonstrate that the ℓ_1 norm of feature gradients is implicitly regularized through the adopted latent perturbation, thereby recovering local linearity and ensuring reliable performance, compared to the existing single-step adversarial training methods. Because latent perturbation is based on the gradients of the latent representations which can be obtained for free in the process of input gradients computation, the proposed method costs roughly the same time as the fast gradient sign method. Experiment results demonstrate that the proposed method, despite its structural simplicity, outperforms state-of-the-art accelerated adversarial training methods.

1. Introduction

Although several studies have suggested the use of deep learning methods to solve challenging tasks, adversarial vulnerability [30] is one of the remaining major challenges while employing deep learning to safety-critical applications. Adversarial training (AT) approaches aim to mitigate the problem by training the model on generated adversarial examples, i.e. the sample corrupted with human-imperceptible noise which can fool the state-of-the-art deep neural networks. Although PGD AT [16] is one of the most effective training methods, it consumes a considerable training time because it relies on multiple projected gradient descent steps to generate the adversaries. The AT based on Fast Gradient

Sign Method (FGSM; [10]) reduces the training time; however, recent works [16, 32, 31] have identified the FGSM’s vulnerability to the sophisticated adversaries.

The trade-off between adversarial robustness and computational cost has facilitated the development of accelerated and trustworthy AT methods. Shafahi et al. [25] significantly reduced the computational burden by presenting a *free* AT method that updates both model parameters and adversarial perturbation through a single shared backward propagation. Wong et al. [38] proposed a *fast* adversarial training based on the discovery that a slight modification in the FGSM training method such as random initialization allows it to achieve an adversarial robustness on par with PGD AT. They also discovered the *catastrophic overfitting* problem of FGSM AT, wherein the model suddenly loses its robustness during training within an epoch.

Although substantial technical advances have been made with regard to the above-mentioned methods, recent works have reported that such approaches are not sufficiently reliable. Andriushchenko et al. [2] demonstrated that fast adversarial training still suffers from the catastrophic overfitting, owing to the deteriorated local linearity of neural networks. Kim et al. [14] found that the fast adversarial training suddenly loses its robustness and eventually collapses when a simple multi-step learning rate schedule is used. Li et al. [15] reported that although fast adversarial training may recover quickly, it still temporally exhibits catastrophic overfitting.

This remaining problem in AT motivates us to explore novel ways to improve the reliability of single-step AT, without bearing a considerable training time. In this study, we demonstrate that the single-step latent adversarial training (SLAT) with the latent adversarial perturbation operates more effectively and reliably compared to the other single-step adversarial training variants. While many of existing adversarial training methods require multiple gradient computations which is inevitably time-consuming, we exploit the gradients of latent representations from multiple layers in *parallel* for the synergistic generation of adversary. Note that the gradients of latent representations can be obtained

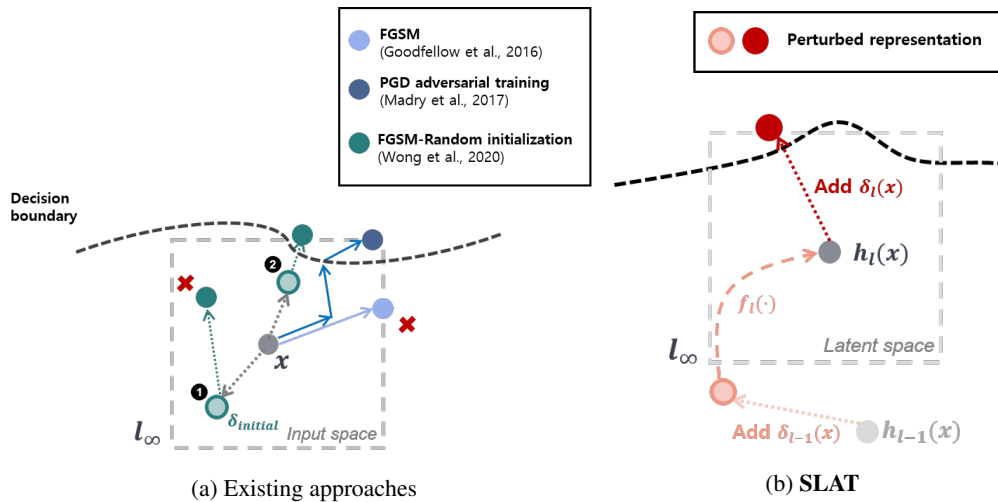


Figure 1: Visual illustration of the conceptual difference between existing and proposed approach. **(a)** FGSM may fail to generate the appropriate adversary because it approximates the solution of inner maximization problem with a single gradient step. While PGD-based AT may generate relatively more desirable adversary, it takes multiple iterations per sample to solve the inner maximization which is computationally expensive. Uniform random initialization [38] contributes to improving the performance of FGSM; however, the success of such initialization is not fundamentally justified. **(b)** Our proposed method mitigates the suggested problems by introducing latent adversarial perturbations in parallel.

for *free* in the process of computation of input gradients.

To the best of our knowledge, our work is the first to reduce the computational cost of AT by deviating from the input-space based frameworks and injecting latent adversarial perturbations. The contribution of this study is summarized as follows. First, we propose that the local linearity of neural networks can be regularized without demanding cost of training time by adopting latent adversarial perturbation, unlike the gradient alignment (GA) regularization [2], which is three times slower compared to FGSM training. In particular, we demonstrate that the ℓ_1 norm of feature gradients is implicitly regularized by introducing latent adversarial perturbation, which closes the gap between the loss function of neural networks and its first-order approximation. As the latent adversarial perturbation is adopted across multiple latent layers, the synergistic regularization effect can be expected. Second, we demonstrate that **SLAT** outperforms the state-of-the-art accelerated adversarial training methods, while achieving performance comparable to PGD AT.

2. Related works

Adversarial training has been improved with the help of adversarial attack algorithms. Goodfellow et al. [10] proposed FGSM to enable rapid generation of adversarial examples through a single-step gradient update. Based on these developments, madry et al. [16] generated a stronger adversary through projected gradient descent (PGD). PGD-based AT has been recognized as a more effective defense method than others, such as provable defenses [37, 42, 6],

label smoothing [24], mix-up [43], and Jacobian regularization [13]. However, the complexity overhead of generating the adversary significantly limits the scalability of the PGD-based AT method.

As introduced in Section 1, several methods [25, 38] have been proposed to improve the efficiency of AT. Zhang et al. [41] analyzed AT from the perspective of a differential game and merged the inner loop of the PGD attack and the gradient update of model parameters. Vivek et al. [34] experimentally found that the adversarial robustness of FGSM AT is improved via application of dropouts to all non-linear layers. Kim et al. [14] demonstrated that the characteristic of FGSM AT, which uses only adversary with the maximum perturbation, leads to the decision boundary distortion and therefore proposed an ad-hoc method to determine an appropriate step size. While most of the proposed methods heavily rely on the input perturbation, we explore the possibility for efficient AT using latent adversarial perturbations.

Although several remarkable works [39, 36] spark interest regarding the importance of latent representations in adversarial robustness, only a few works have directly leveraged the latent adversarial perturbations for AT. Sankaranarayanan et al. [23] used the gradients of latent representations computed from the preceding mini-batch to approximate the solution of inner-maximization. Although the proposed method has yielded modest improvements in terms of the adversarial robustness with respect to FGSM-based AT, they do not necessarily mean that a latent adversarial perturbation from a gradient of the preceding mini-batch is optimal. While [27]

suggested latent adversarial training for further fine-tuning of the adversarially trained model, it is contingent on performance of multi-step PGD AT. On the contrary, our approach focusses on the reduction of computational costs of AT via latent adversarial perturbations.

3. Latent adversarial perturbation

We begin with formulating the generalized min-max adversarial training objective in terms of latent adversarial perturbation. Let $(x \in \mathcal{X}, y \in \mathcal{Y}) \sim D$ be the pair of sample and label instance generated from the distribution D , given sample space \mathcal{X} and label space \mathcal{Y} . We represent $f_l(\cdot)$ as the function defined by the l -th layer, and $h_l(x)$ as the latent-representation vector given sample x , where $h_0(x) \triangleq x$. Precisely, $h_l(x) = f_l(f_{l-1}(\dots(f_1(x)))) = f_l(h_{l-1}(x))$, where $f_l(h_{l-1}(x)) = \phi(W_l h_{l-1}(x) + b_l)$, given the l -th weight matrix W_l , bias vector b_l , and the activation function $\phi(\cdot)$. Note that the latent adversarial perturbation is not considered yet. We denote the L -layer neural networks as a function $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, parameterized by $\theta = \{W_1, \dots, W_L, b_1, \dots, b_L\}$: $f_\theta(x) = f_L(f_{L-1}(\dots(f_1(x))))$. For simplicity, let $f_{m:n}(x) = f_n(f_{n-1}(\dots(f_m(h_{m-1}(x)))))$, for any $1 \leq m < n \leq L$.

In this study, we investigate the benefits of latent adversarial perturbation, which is built based on the gradients of latent representations. Let $K \subseteq \{0, \dots, L-1\}$ be the subset of layer indexes injected with adversarial perturbation, where the 0th layer represents the input layer. We denote the adversarial perturbation given x for the k -th layer as $\delta_k(x)$, where $k \in K$, and the set of adversarial perturbations is $\delta = \{\delta_k(x)\}_{\forall k \in K}$.

To examine the marginalized effect of latent adversarial perturbations, we define the *accumulated* perturbation in layer $L-1$ as $\hat{\delta}_{L-1}(x)$, where $\hat{\delta}_{L-1}(x)$ originates from the forward propagation of δ . Note that we virtually introduced the accumulated perturbation for the sake of the analysis. In practice, each latent adversarial perturbation is applied layer-wise. Thus, $h_{l+1}(x) = f_l(h_l(x) + \delta_l(x))$, $\forall l \in K$.

The optimal set of parameters θ^* and adversarial perturbations δ^* can be obtained by solving the following min-max problem:

$$\theta^*, \delta^* = \arg \min_{\theta} \mathbb{E}_{(x,y) \sim D} \max_{\delta} \left[\mathcal{L} \left(f_L(h_{L-1}(x) + \hat{\delta}_{L-1}(x)), y \right) \right]. \quad (1)$$

While it is difficult to obtain the accumulated perturbation $\hat{\delta}_{L-1}(x)$ in a closed form given highly nonlinear function $f_\theta(\cdot)$, we can approximate it by making a reasonable assumption that $\delta_k(x)$ is sufficiently small. The approximation is based on [4] which examined the effect of Gaussian Noise Injection (GNI) into multiple latent layers. The adversarial

perturbation accumulated on the layer $L-1$ can be expressed as follows:

Proposition 1. Consider an L layer neural network, with the latent adversarial perturbations $\delta_k(x)$ being applied at each layer $k \in K$. Assuming the Hessians, of the form $\nabla^2 h_l(x)|_{h_m(x)}$ where l, m are the index over layers, are finite. Then the perturbation accumulated at the layer $L-1$, $\hat{\delta}_{L-1}(x)$, is approximated by:

$$\hat{\delta}_{L-1}(x) = \sum_{k \in K} \mathbf{J}_k(x) \delta_k(x) + O(\gamma), \quad (2)$$

where $\mathbf{J}_k(x) \in \mathbb{R}^{N_{L-1} \times N_k}$ represents each layer's Jacobian; $\mathbf{J}_k(x)_{i,j} = \frac{\partial h_{L-1}(x)_i}{\partial h_k(x)_j}$, given the number of neurons in layer $L-1$ and k as N_{L-1} and N_k , respectively. $O(\gamma)$ represents higher order terms in δ that tend to zero in the limit of small perturbation.

The detailed proof is provided in the supplementary material. Based on the framework (1) and Proposition 1, we provide the details of the proposed latent adversarial perturbation. By neglecting the higher order terms in Proposition 1, the linear approximation of the loss function $\mathcal{L}(f_L(h_{L-1}(x) + \hat{\delta}_{L-1}(x)), y)$ is as follows:

$$\begin{aligned} & \mathcal{L} \left(f_L(h_{L-1}(x) + \hat{\delta}_{L-1}(x)), y \right) \\ & \approx \mathcal{L} \left(f_L(h_{L-1}(x)), y \right) + \\ & \quad \nabla_{h_{L-1}(x)} \mathcal{L} \left(f_L(h_{L-1}(x)), y \right)^T \hat{\delta}_{L-1}(x) \\ & = \mathcal{L} \left(f_L(h_{L-1}(x)), y \right) + \\ & \quad \nabla_{h_{L-1}(x)} \mathcal{L} \left(f_L(h_{L-1}(x)), y \right)^T \sum_{k \in K} \mathbf{J}_k(x) \delta_k(x). \end{aligned} \quad (3)$$

Then, we approximate the solution of the inner maximization problem in (1), as similarly done in FGSM:

$$\begin{aligned} \delta_k(x) & = \eta_k \cdot \text{sign} \left(\mathbf{J}_k(x)^T \nabla_{h_{L-1}(x)} \mathcal{L} \left(f_L(h_L(x)), y \right) \right) \\ & = \eta_k \cdot \text{sign} \left(\nabla_{h_k(x)} \mathcal{L} \left(f_\theta(x), y \right) \right), \forall k \in K, \end{aligned} \quad (4)$$

where η_k is the step size for the k -th layer. Then, the explicit regularizer of latent adversarial perturbation in (3) is derived as follows:

$$\begin{aligned} \mathcal{L}_{\text{latent}} & = \sum_{k \in K} \eta_k \cdot \text{sign} \left(\nabla_{h_k(x)} \mathcal{L} \left(f_\theta(x), y \right) \right) \\ & \quad \circ \nabla_{h_k(x)} \mathcal{L} \left(f_\theta(x), y \right) \\ & = \sum_{k \in K} \eta_k \cdot \|\nabla_{h_k(x)} \mathcal{L} \left(f_\theta(x), y \right)\|_1, \end{aligned} \quad (5)$$

where \circ represents dot product.

It affords us the theoretical insights that the latent adversarial perturbation leads to the implicit regularization of ℓ_1 norm of the feature gradients. Although the uses of input gradient regularization in adversarial defense [21] or explainable machine learning [22, 28] have been widely recognized, the effects of feature gradient regularization on adversarial robustness have been poorly understood. To address this, we establish a connection between the feature gradient regularization and local linearity. For a better linear approximation, the linear approximation error $R_k(x)$ should be constrained:

$$\left| \mathcal{L}\left(f_{k:L}(h_{k-1}(x) + \epsilon), y\right) - \mathcal{L}\left(f_{k:L}(h_{k-1}(x)), y\right) - \left\langle \nabla_{h_{k-1}(x)} \mathcal{L}\left(f_{k:L}(h_{k-1}(x)), y\right), \epsilon \right\rangle \right|, \quad (6)$$

where ϵ is an arbitrary perturbation with sufficiently small size, and $R_k(x)$ is defined with the function $f_{k:L}(\cdot)$ for an arbitrary $k \in K$. Note that $R_k(x)$ includes the second-order term $\langle \epsilon, H_k(x)\epsilon \rangle$ where $H_k(x) = \nabla_{h_{k-1}(x)}^2 \mathcal{L}(f_{k:L}(h_{k-1}(x)), y)$. Let ∇_{k-1} be a shorthand for $\nabla_{h_{k-1}(x)} \mathcal{L}(f_{k:L}(h_{k-1}(x)), y)$. By the low-rank approximation of the Hessian matrix [17], let $H_k(x) \approx \nabla_{k-1} \nabla_{k-1}^T$. Then, the upper bound of the second-order term is as follows:

$$\begin{aligned} \langle \epsilon, H_k(x)\epsilon \rangle &\approx |\langle \epsilon, \nabla_{k-1} \rangle|^2 \\ &\leq \|\epsilon\|_2^2 \|\nabla_{k-1}\|_2^2 \\ &\leq \|\epsilon\|_2^2 \|\nabla_{k-1}\|_1^2, \end{aligned} \quad (7)$$

since $\|x\|_p \geq \|x\|_q$ for $0 < p < q, \forall x \in \mathbb{R}^n$. Thus the regularization of ℓ_1 norm of feature gradients may result in a better linear approximation of the loss function. This eventually contributes to improving the reliability of FGSM which relies heavily on linear approximation of the loss function.

We further foster a close collaboration between feature gradient regularization and the minimization of adversarial loss. Inspired from [26], the small variation in the loss $\Delta \mathcal{L}_k$ caused by the latent adversarial perturbation $\delta_k(x)$ is as follows:

$$\begin{aligned} \Delta \mathcal{L}_k &= \max_{\|\delta_k(x)\| \leq \eta_k} \left| \mathcal{L}\left(f_{k+1:L}(h_k(x) + \delta_k(x)), y\right) - \mathcal{L}\left(f_{k+1:L}(h_k(x)), y\right) \right| \\ &\approx \max_{\|\delta_k(x)\| \leq \eta_k} \left| \left\langle \nabla_{h_k(x)} \mathcal{L}\left(f_{k+1:L}(h_k(x)), y\right), \delta_k(x) \right\rangle \right| \\ &= \eta_k \left\| \nabla_{h_k(x)} \mathcal{L}\left(f_{k+1:L}(h_k(x)), y\right) \right\|_*, \end{aligned} \quad (8)$$

where η_k is the allowed step size for the perturbation $\delta_k(x)$. The last equality comes from the definition of the dual norm $\|\cdot\|_*$ of $\|\cdot\|$. Thus the regularization of ℓ_1 norm of feature gradients is closely related to minimizing the adversarial loss stems from $\delta_k(x)$ with limited ℓ_∞ norm.

Advantages over other local linearity regularization methods. Existing works have been proposed to theoretically investigate the underlying principles of adversarial training with respect to local linearity. Moosavi et al. [18] demonstrated that adversarial training increases adversarial robustness by decreasing curvature and proposed a new curvature regularizer based on the finite difference approximation of Hessian. While the proposed regularizer improves the adversarial robustness with affordable computational costs, it requires pretrained networks for fine-tuning. Moreover, the reported standard and adversarial accuracy was lower than that of our proposed method given the same base network architecture (Section 4).

A few follow-up studies have attempted to regularize linear approximation errors in an optimal manner. Qin et al. [20] explores the worst-case perturbation that maximizes the linear approximation error of the loss function via multiple steps of PGD in order to regularize local linearity and penalize gradient obfuscation [3]. Tsiligkaridis et al. [33] proposed a novel adversarial training method based on the Frank-Wolfe algorithm [12] to decrease the directional variation of loss gradients. Compared to existing works that inevitably rely on the multi-step gradient updates for regularizing local linearity, the above latent adversarial perturbation (4) enables us to seamlessly regularize the local linearity of sub-networks in parallel with affordable computational costs.

4. Experiments

4.1. Experiments on Toy Dataset

To conceptually clarify the effects of latent adversarial perturbation, we observed the behavior of the classifier on a simple binary classification problem. We generate samples (x, y) from 2D gaussian distributions, $\mathcal{N}_1(\mu, \Sigma)$, and $\mathcal{N}_2(-\mu, \Sigma)$. Ilyas et al. [11] proposed that the adversarial samples stem from useful but non-robust features which are informative for classification in standard setting but vulnerable to adversarial attacks. Inspired by [11], we intentionally compose the input features as robust (on X-axis) and non-robust feature (on Y-axis, Figure 2) in order to investigate whether the latent adversarial perturbation helps filtering out the non-robust feature.

A simple neural network ($L = 2$) is implemented where the adversarial perturbation is injected to each input layer and latent layer, i.e., $K = \{0, 1\}$. Then we compare the decision boundary of binary classifiers obtained through standard training, FGSM AT, and SLAT ($\eta_0 = 0.1$ for both

Algorithm 1: Single-step Latent Adversarial Training method (SLAT)

Input: Training iteration T , Number of samples N , Number of layers L , Training set $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$, Subset of layer indexes K , Layer-wise step size η_k

Output: Adversarially robust network f_θ

```
for  $t \leftarrow 1$  to  $T$  do
  for  $i \leftarrow 1$  to  $N$  do
    for  $k \in K$  do
      // Compute latent adversarial perturbations
       $\delta_k(x_i) = \eta_k \cdot \text{sign}(\nabla_{h_k(x_i)} \mathcal{L}(f_\theta(x_i), y_i))$ 
    for  $l \in \{0, \dots, L-2\}$  do
      if  $l \in K$  then
        // Propagate adversarial perturbations forward
         $h_{l+1}(x_i) = f_{l+1}(h_l(x_i) + \delta_l(x_i))$ 
      else
         $h_{l+1}(x_i) = f_{l+1}(h_l(x_i))$ 
    Optimize  $\theta$  by the objective  $\mathcal{L}(f_L(h_{L-1}(x_i)), y_i)$  using gradient descent.
```

FGSM AT and SLAT). Under the limited adversarial budgets, while FGSM AT still heavily relies on the predictive but non-robust feature (Figure 2b), SLAT learns to select a more robust feature (Figure 2c). It implies that the latent adversarial perturbation collaborates with the input adversarial perturbation, so that the effect of adversarial training can be amplified.

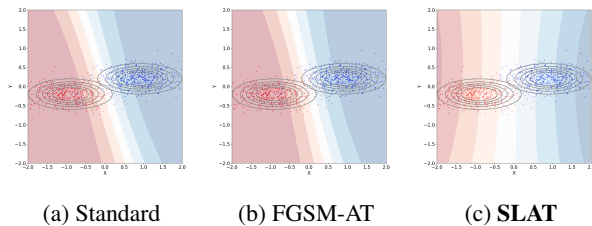


Figure 2: Decision boundary of binary classifiers. Best viewed in color.

4.2. Comparison of robustness

Datasets. To further investigate the effect of latent adversarial perturbation on adversarial robustness, we compare the adversarial robustness of several models on CIFAR-10, CIFAR-100 and Tiny ImageNet. For CIFAR-10 and CIFAR-100, all the images are randomly cropped into 32×32 with

padding size 4 following [38, 2]. Images in Tiny ImageNet are similarly cropped into 64×64 .

Experimental setup. For an adversarial example, ℓ_∞ -perturbation is used with radius $\eta_0 = 8/255$. For a fair comparison, we reproduced all the other baseline results using the same back-bone architecture and the optimization settings. Specifically, we validate several AT methods using Wide ResNet 28-10 [40], and cyclic learning rates [29] with the SGD optimizer following the setup of [2].

The latent adversarial perturbation is injected into three layers (K), including input layer and last layers in each group *conv1, conv2* [40]. We use $\eta_k = 8/255$ for every $k \in K$. We evaluate the adversarial robustness of several models using PGD-50-10 attack [16], i.e., with 50 iterations and 10 restarts, and AutoAttack [8] which is the ensemble of two extended PGD attacks, a white-box FAB-attack [7], and the black-box Square Attack [1], for verifying the absence of gradient masking [3]. Every experiments are run on a single GeForce Titan X. The details regarding the simulation settings are presented in the supplementary material.

Baselines. We compare our method with the following state-of-the-art (fast) adversarial training methods: PGD AT [16], FGSM-GA (FGSM AT with gradient alignment, [2]), YOPO [41], Free-AT [25], FGSM AT [10], FGSM-RS (FGSM AT with random step, [38]), and FGSM-CKPT (FGSM AT with checkpoints, [14]). Every model is trained based on the reported hyperparameter configurations of original papers except FGSM-GA which experiences catastrophic overfitting with the reported gradient alignment parameter $\lambda = 0.2$. We therefore tested with the parameter $\lambda \in \{0.2, 0.3, 0.5, 1.0, 2.0\}$ where $\lambda = 1.0$ is selected.

Results. The clean and robust accuracy for the CIFAR-10, CIFAR-100 and Tiny ImageNet datasets are summarized in Table 1. We found that SLAT reliably outperforms most of the accelerated adversarial training methods, including YOPO [41], Free-AT [25], FGSM-RS, and FGSM-CKPT [14], with respect to robust accuracy against PGD attack and AutoAttack [8], on every experiments. Note that most of the FGSM-based accelerated AT methods experienced a catastrophic overfitting during the training process on specific datasets, e.g., FGSM on CIFAR-10, CIFAR-100; FGSM-CKPT on CIFAR-100, Tiny ImageNet; FGSM-RS on CIFAR-10, CIFAR-100. Although FGSM-GA [2] and PGD-7 tend to exhibit better performance than SLAT, both methods are much slower than the other adversarial training methods. Moreover, SLAT achieves a higher standard accuracy than both methods on CIFAR-10.

4.3. Quantitative and Qualitative analysis

Visualizing loss landscape. To fully justify the association between latent hidden perturbation and local linearity, we visualize the loss landscape of several models. Figure 3d shows that the loss of SLAT in ℓ_∞ ball is almost perfectly in

Table 1: Standard and robust accuracies (%) on CIFAR-10, CIFAR-100, and Tiny ImageNet datasets.

	Method	Standard	PGD-50-10	AutoAttack	Training time (min)
CIFAR-10	PGD-7	84.86±0.16	51.63±0.13	48.65±0.08	383.2
	FGSM-GA	82.88±0.01	48.90±0.37	46.22±0.30	297.9
	YOPO-5-3	82.35±1.78	34.23±3.61	32.79±3.65	62.5
	Free-AT ($m = 8$)	76.57±0.19	44.15±0.30	41.02±0.20	119.4
	FGSM	87.42±1.08	0.01±0.01	0.00±0.00	100.5
	FGSM-RS	90.76±6.36	3.90±4.06	0.44±0.50	99.7
	FGSM-CKPT ($c = 3$)	89.32±0.10	40.83±0.36	39.38±0.24	121.4
SLAT	85.91±0.31	47.06±0.03	44.62±0.11	104.6	
CIFAR-100	PGD-7	59.59±0.17	29.58±0.24	26.00±0.20	392.1
	FGSM-GA	58.63±0.17	27.53±0.10	24.07±0.15	240.5
	YOPO-5-3	51.45±7.33	15.23±2.01	13.94±1.82	65.0
	Free-AT ($m = 8$)	48.02±0.29	22.40±0.19	18.67±0.03	117.1
	FGSM	61.96±2.17	0.00±0.00	0.00±0.00	99.9
	FGSM-RS	50.96±4.57	0.00±0.00	0.00±0.00	100.9
	FGSM-CKPT ($c = 3$)	73.53±0.65	0.66±0.60	0.09±0.09	101.5
SLAT	59.56±0.50	26.26±0.47	23.02±0.14	101.7	
Tiny ImageNet	PGD-7	48.92±0.43	23.05±0.35	18.78±0.14	3098.3
	FGSM-GA	48.73±0.14	22.62±0.11	18.34±0.07	2032.2
	YOPO-5-3	51.45±6.01	15.08±1.78	13.94±1.61	511.5
	Free-AT ($m = 8$)	22.40±0.17	9.05±0.08	6.06±0.18	911.6
	FGSM	36.47±11.75	8.68±12.27	6.63±9.38	779.8
	FGSM-RS	42.13±14.98	10.32±11.93	8.41±9.73	787.5
	FGSM-CKPT ($c = 3$)	61.64±2.24	5.91±6.68	5.26±5.98	753.0
SLAT	48.77±0.25	20.21±0.16	16.38±0.16	785.5	

linear with the adversarial direction, thus qualitatively proving that the local linearity is recovered by adopting latent perturbation. Furthermore, Figure 3a, 3b shows that the loss landscape of FGSM and FGSM-RS are highly non-linear in the vicinity of training examples. Note that the adversarial loss of FGSM-RS given adversary with maximum step size is relatively lower than the other well-trained methods (Figure 3c, 3d). Due to the catastrophic overfitting problem, both methods are still vulnerable to other smaller perturbations. Although FGSM-CKPT [14] performs better than FGSM-RS, the obtained loss landscape is not perfectly linear as SLAT in adversarial direction.

Evaluating adversarial robustness with different ℓ_∞ radius. To investigate the contribution of latent adversarial perturbation on reliability, we compare the robustness of FGSM-based adversarial training methods with different ℓ_∞ -radius (Figure 4a). The results illustrate that the latent adversarial perturbation prevents the model from losing its robustness quickly when the ℓ_∞ -radius increases. Note that the FGSM-based single-step adversarial training methods (FGSM, FGSM-RS, FGSM-CKPT) are relatively sensitive to the ℓ_∞ -radius, potentially due to the lack of regularization

for a better linear approximation.

Measuring local linearity. To specify the connection between catastrophic over-fitting and local linearity, we measure the adversarial robustness and gradient alignment as done in [1] (Figure 4b, 4c, 4d). Gradient alignment is measured as the cosine similarity between input gradients computed with original and randomly perturbed sample. We figured out that the gradient alignment suddenly drops as the model loses its adversarial robustness. While FGSM or FGSM-RS may rely on the early-stopping technique [38] to empirically prevent catastrophic overfitting, it is not sufficient to outperform SLAT (Figure 4b).

Measuring ℓ_1 norm of feature gradients. With respect to the analysis in the section 3, we compared the ℓ_1 norm of the stochastic gradients computed with different methods. After training, we computed the ℓ_1 norm of gradients for every representation $h_l(x)$ for all $l \in K$. Figure 5 shows that SLAT yields lower norm of the gradients compared to the other methods on layers injected with latent perturbation (*Input, Conv1, Conv2*). It is potentially due to the synergistic regularization by latent adversarial perturbation.

Ablation study on latent adversarial perturbation. To

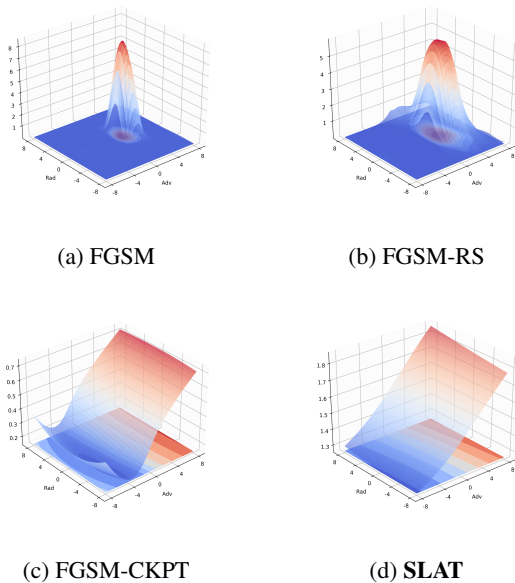


Figure 3: Visualization of loss landscape on CIFAR-10 for various models. We plot the softmax cross entropy loss projected on adversarial direction and random (Rademacher) direction with $\eta_0 = 8/255$ radius.

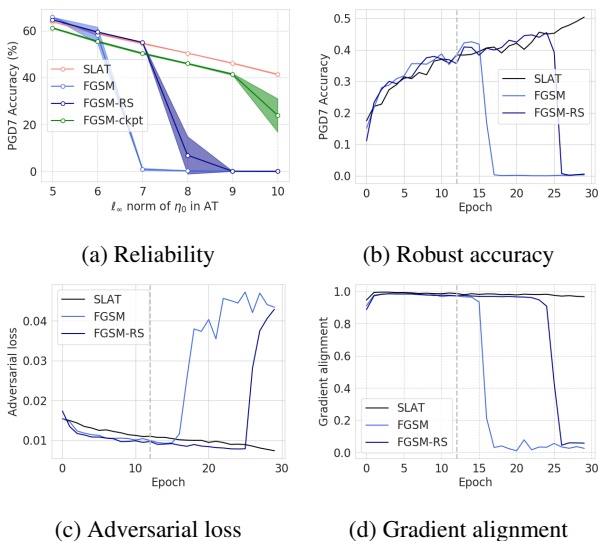


Figure 4: **(a)** Robust accuracy (%) of various AT methods with different l_∞ radius on CIFAR-10. The results are averaged on 4 different random seeds. **(b)~(d)** Demonstrating the relationship between catastrophic overfitting and local linearity on CIFAR-10. Gray dotted line indicates the epoch at maximum learning rate.

quantify the extent of performance improvement achieved by the latent adversarial perturbation, we compared the standard and robust accuracy of the different versions of **SLAT** (Ta-

ble 2). When the latent adversarial perturbation is excluded, the proposed method reduces to the conventional FGSM AT. To consider the effect of step size reduction [2], we also measured the performance of the version of FGSM AT with a reduced step size. Table 2 shows that the vanilla FGSM AT is not sufficient to prevent the catastrophic overfitting problem regardless of step size. Moreover, we found that latent adversarial perturbations prevent the catastrophic overfitting of FGSM-RS [38]. This result indicates that the latent adversarial perturbation plays an important role in enhancing the reliability of adversarial training. Although the latent adversarial perturbations improve the adversarial robustness of FGSM-RS to some extent, it does not outperform **SLAT**. It is because the latent adversarial perturbation was derived from the original sample rather than the randomly perturbed sample as in (4).

Table 2: Ablation study on latent adversarial perturbation with $\eta_0 = 8/255$ (CIFAR-10).

Method	Standard	PGD-50-10
FGSM (step size= $8/255$)	87.42	0.01
FGSM (step size= $0.9 * 8/255$)	90.87	3.09
FGSM-RS (w/ latent perturbation)	82.87	44.95
SLAT	85.91	47.06

Moreover, to understand the effect of layer depth on latent adversarial training, we compared the performance of **SLAT** with different subsets of layer indexes K . Following [40], latent adversarial perturbation was added to the last layer of some selected blocks among *conv1*, *conv2*, *conv3*, *conv4* of Wide ResNet 28-10. The input layer was included equally in K for all experiments. We experimentally found that the robustness and standard accuracy decreased significantly when the latent adversarial perturbations were injected into deeper layers (Table 3). This implies that simply injecting latent adversarial perturbation to arbitrarily many hidden layers is not necessarily effective and may over-regularize the networks. These findings provide us an interesting conjecture that the recovery of local linearity should be primarily confined to early sub-networks that recursively include other deeper sub-networks. The theoretical analysis is left for future research.

We additionally conducted the analysis on hyperparameter sensitivity with the adversarial step size η_k (Figure 6a). Standard and adversarial accuracy are measured on CIFAR-10, where the step size η_k varies from $0.6 * 8/255$ to $1.0 * 8/255$. The robust accuracy is high when $\eta_k = 8/255$.

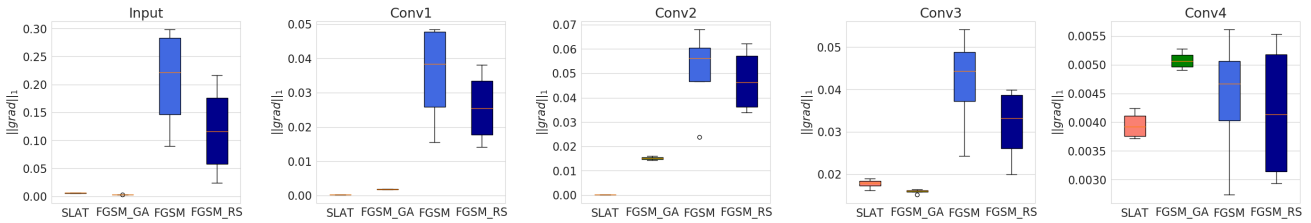


Figure 5: ℓ_1 norm of gradients on CIFAR-10 after training. The results are averaged on 4 different random seeds.

Table 3: Analysis of adversarial robustness and standard accuracy based on layer depth change (CIFAR-10).

Layers	Standard	PGD-50-10	AutoAttack
<i>conv1, conv2</i> (SLAT)	85.91	47.06	44.62
<i>conv3, conv4</i>	84.60	0.00	0.00
<i>conv1, conv2,</i> <i>conv3, conv4</i>	81.41	47.3	43.58

logits of FGSM and R+FGSM [31] adversary to confirm the absence of gradient masking following [5, 35, 31]. If the model exhibits gradient masking, the distance between logits of FGSM and R+FGSM increases due to the sharp curvature of loss landscape in vicinity of training examples. Figure 6b shows that the distance of proposed method is relatively low compared to FGSM AT during the training process, thus proving that the proposed method does not rely on gradient masking.

5. Conclusion

In this study, we demonstrate that the latent adversarial perturbation may provide a novel breakthrough for the efficient AT. The proposed framework allows us to compensate for the local linearity without sacrificing training time. Further, we establish a bridge between latent adversarial perturbation and adversarial loss minimization. It enables us to learn adversarially robust model in a more reliable manner, compared to the recent fast adversarial training method [38] which lacks any form of regularization. By running simulations on various benchmark datasets, we illustrate that our model significantly outperforms state-of-the-art accelerated adversarial training methods. The proposed method is fully-architecture agnostic, has only a few free parameters to tune, and is potentially compatible with many other AT methods. We believe that expanding the proposed framework beyond single-step AT will be an interesting future work.

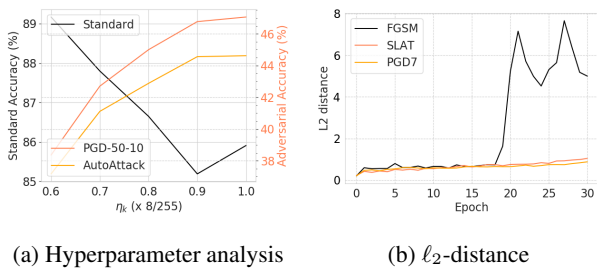


Figure 6: (a) Impact of η_k on robustness. (b) Average ℓ_2 distance between logits of FGSM and R+FGSM [31] adversary (CIFAR-10).

Thus, we fix η_k as $8/255$ for every dataset and layer $k \in K$. It may work better to fine-tune η_k differently depending on k . From our preliminary analyses, the unified η_k for the latent representations worked sufficiently well that we did not feel the need to fine-tune η_k for every layer. Moreover, while the trade-off between robustness and accuracy is observed, the standard accuracy for $\eta_k = 8/255$ is still superior than that of FGSM-GA or PGD-7 (Table 1).

Verifying absence of gradient masking. Although many certified adversarial defense methods have reported substantial advances in adversarial robustness, recent works [19, 9, 3] report that the gradient masking may present a false sense of security, i.e., making it difficult to generate adversary using gradient methods by obfuscating gradients, while the adversary still potentially exists. Besides the quantitative (Autoattack in Table 1) and qualitative (Figure 3, [20, 25]) evidences, we additionally measure the ℓ_2 distance between

Acknowledgments

This work was supported by Institute for Information & Communications Technology Promotion (IITP) grant funded by the Korea government (No. 2017-0-00451) (No.2019-0-01371, Development of brain-inspired AI with human-like intelligence), National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (NRF-2019M3E5D2A01066267), and Samsung Research Funding Center of Samsung Electronics under Project Number SRFC-TC1603-06.

References

- [1] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: a query-efficient black-box adversarial attack via random search. In *European Conference on Computer Vision*, pages 484–501. Springer, 2020. 5, 6
- [2] Maksym Andriushchenko and Nicolas Flammarion. Understanding and improving fast adversarial training. *Advances in Neural Information Processing Systems*, 33, 2020. 1, 2, 5, 7
- [3] 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*, pages 274–283. PMLR, 2018. 4, 5, 8
- [4] Alexander Camuto, Matthew Willetts, Umut Şimşekli, Stephen Roberts, and Chris Holmes. Explicit regularisation in gaussian noise injections. *arXiv preprint arXiv:2007.07368*, 2020. 3
- [5] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, Aleksander Madry, and Alexey Kurakin. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019. 8
- [6] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, pages 1310–1320. PMLR, 2019. 2
- [7] Francesco Croce and Matthias Hein. Minimally distorted adversarial examples with a fast adaptive boundary attack. In *International Conference on Machine Learning*, pages 2196–2205. PMLR, 2020. 5
- [8] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International Conference on Machine Learning*, pages 2206–2216. PMLR, 2020. 5
- [9] Logan Engstrom, Andrew Ilyas, and Anish Athalye. Evaluating and understanding the robustness of adversarial logit pairing. *arXiv preprint arXiv:1807.10272*, 2018. 8
- [10] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 1, 2, 5
- [11] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. In *Advances in Neural Information Processing Systems*, pages 125–136, 2019. 4
- [12] Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th international conference on machine learning*, number CONF, pages 427–435, 2013. 4
- [13] Daniel Jakubovitz and Raja Giryes. Improving dnn robustness to adversarial attacks using jacobian regularization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 514–529, 2018. 2
- [14] Hoki Kim, Woojin Lee, and Jaewook Lee. Understanding catastrophic overfitting in single-step adversarial training. *arXiv preprint arXiv:2010.01799*, 2020. 1, 2, 5, 6
- [15] Bai Li, Shiqi Wang, Suman Jana, and Lawrence Carin. Towards understanding fast adversarial training. *arXiv preprint arXiv:2006.03089*, 2020. 1
- [16] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. 1, 2, 5
- [17] James Martens, Ilya Sutskever, and Kevin Swersky. Estimating the hessian by back-propagating curvature. *arXiv preprint arXiv:1206.6464*, 2012. 4
- [18] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Jonathan Uesato, and Pascal Frossard. Robustness via curvature regularization, and vice versa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9078–9086, 2019. 4
- [19] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, pages 506–519, 2017. 8
- [20] Chongli Qin, James Martens, Sven Gowal, Dilip Krishnan, Krishnamurthy Dvijotham, Alhussein Fawzi, Soham De, Robert Stanforth, and Pushmeet Kohli. Adversarial robustness through local linearization. *arXiv preprint arXiv:1907.02610*, 2019. 4, 8
- [21] Andrew Ross and Finale Doshi-Velez. Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 4
- [22] Andrew Ross, Isaac Lage, and Finale Doshi-Velez. The neural lasso: Local linear sparsity for interpretable explanations. In *Workshop on Transparent and Interpretable Machine Learning in Safety Critical Environments, 31st Conference on Neural Information Processing Systems*, 2017. 4
- [23] Swami Sankaranarayanan, Arpit Jain, Rama Chellappa, and Ser Nam Lim. Regularizing deep networks using efficient layerwise adversarial training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 2
- [24] Ali Shafahi, Amin Ghiasi, Furong Huang, and Tom Goldstein. Label smoothing and logit squeezing: a replacement for adversarial training? *arXiv preprint arXiv:1910.11585*, 2019. 2
- [25] Ali Shafahi, Mahyar Najibi, Mohammad Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! In *Advances in Neural Information Processing Systems*, pages 3358–3369, 2019. 1, 2, 5, 8
- [26] Carl-Johann Simon-Gabriel, Yann Ollivier, Leon Bottou, Bernhard Schölkopf, and David Lopez-Paz. First-order adversarial vulnerability of neural networks and input dimension. In *International Conference on Machine Learning*, pages 5809–5817. PMLR, 2019. 4
- [27] Mayank Singh, Abhishek Sinha, Nupur Kumari, Harshitha Machiraju, Balaji Krishnamurthy, and Vineeth N Balasubramanian. Harnessing the vulnerability of latent layers in adversarially trained models. *arXiv preprint arXiv:1905.05186*, 2019. 2

- [28] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017. 4
- [29] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, volume 11006, page 1100612. International Society for Optics and Photonics, 2019. 5
- [30] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. 1
- [31] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*, 2017. 1, 8
- [32] Florian Tramèr, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. The space of transferable adversarial examples. *arXiv preprint arXiv:1704.03453*, 2017. 1
- [33] Theodoros Tsiligkaridis and Jay Roberts. On frank-wolfe optimization for adversarial robustness and interpretability. *arXiv preprint arXiv:2012.12368*, 2020. 4
- [34] BS Vivek and R Venkatesh Babu. Single-step adversarial training with dropout scheduling. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 947–956. IEEE, 2020. 2
- [35] BS Vivek, Arya Baburaj, and R Venkatesh Babu. Regularizer to mitigate gradient masking effect during single-step adversarial training. In *CVPR Workshops*, pages 66–73, 2019. 8
- [36] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P Xing. High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8684–8694, 2020. 2
- [37] Eric Wong and Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning*, pages 5286–5295. PMLR, 2018. 2
- [38] Eric Wong, Leslie Rice, and J Zico Kolter. Fast is better than free: Revisiting adversarial training. *arXiv preprint arXiv:2001.03994*, 2020. 1, 2, 5, 6, 7, 8
- [39] Cihang Xie, Yuxin Wu, Laurens van der Maaten, Alan L Yuille, and Kaiming He. Feature denoising for improving adversarial robustness. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 501–509, 2019. 2
- [40] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. 5, 7
- [41] Dinghui Zhang, Tianyuan Zhang, Yiping Lu, Zhanxing Zhu, and Bin Dong. You only propagate once: Accelerating adversarial training via maximal principle. *arXiv preprint arXiv:1905.00877*, 2019. 2, 5
- [42] Huan Zhang, Hongge Chen, Chaowei Xiao, Sven Gowal, Robert Stanforth, Bo Li, Duane Boning, and Cho-Jui Hsieh. Towards stable and efficient training of verifiably robust neural networks. *arXiv preprint arXiv:1906.06316*, 2019. 2
- [43] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 2