

# Randomized Adversarial Training via Taylor Expansion

Gaojie Jin<sup>1,2</sup>, Xinping Yi<sup>2</sup>, Dengyu Wu<sup>2</sup>, Ronghui Mu<sup>3</sup>, Xiaowei Huang<sup>2</sup>

<sup>1</sup>State Key Laboratory of Computer Science, Institute of Software, CAS, Beijing, China

<sup>2</sup>University of Liverpool, Liverpool, UK

<sup>3</sup>Lancaster University, Lancaster, UK

{g.jin3, xinping.yi, xiaowei.huang}@liverpool.ac.uk

 Corresponding Author

## Abstract

*In recent years, there has been an explosion of research into developing more robust deep neural networks against adversarial examples. Adversarial training appears as one of the most successful methods. To deal with both the robustness against adversarial examples and the accuracy over clean examples, many works develop enhanced adversarial training methods to achieve various trade-offs between them [19, 38, 80]. Leveraging over the studies [8, 32] that smoothed update on weights during training may help find flat minima and improve generalization, we suggest reconciling the robustness-accuracy trade-off from another perspective, i.e., by adding random noise into deterministic weights. The randomized weights enable our design of a novel adversarial training method via Taylor expansion of a small Gaussian noise, and we show that the new adversarial training method can flatten loss landscape and find flat minima. With PGD, CW, and Auto Attacks, an extensive set of experiments demonstrate that our method enhances the state-of-the-art adversarial training methods, boosting both robustness and clean accuracy. The code is available at <https://github.com/Alexkael/Randomized-Adversarial-Training>.*

## 1. Introduction

Trade-off between adversarial robustness and clean accuracy has recently been intensively studied [63, 68, 80] and demonstrated to exist [34, 58, 70, 77]. Many different techniques have been developed to alleviate the loss of clean accuracy when improving robustness, including data augmentation [1, 5, 29], early-stopping [60, 81], instance reweighting [3, 82], and various adversarial training methods [11, 36, 42, 74, 80]. Adversarial training is believed to be the most effective defense method against adversarial at-

tacks, and is usually formulated as a minimax optimization problem where the network weights are assumed deterministic in each alternating iteration. Given the fact that both clean and adversarial examples are drawn from unknown distributions which interact with one another through the network weights, it is reasonable to relax the assumption that neural networks are *simply* deterministic models where weights are scalar values. This paper is based on a view, illustrated in Fig. 1, that randomized models enable the training optimization to consider multiple directions within a small area and may achieve smoothed weights update – in a way different from checkpoint averaging [8, 32] – and obtain robust models against new clean/adversarial examples.

Building upon the above view, we find a way, drastically different from most existing studies in adversarial training, to balance robustness and clean accuracy, that is, *embedding neural network weights with random noise*. Whilst the randomized weights framework is not new in statistical learning theory, where it has been used in many previous works for e.g., generalization analysis [18, 52, 75], we hope to advance the empirical understanding of the robustness-accuracy trade-off problem in adversarial training by leveraging the rich tool sets in statistical learning. Remarkably, it turns out adversarial training with the optimization over randomized weights can improve the state-of-the-art adversarial training methods over both the adversarial robustness and clean accuracy.

By modeling weights as randomized variables with an artificially injected weight perturbation, we start with an empirical analysis of flatness of loss landscape in Sec. 3. We show that our method can flatten loss landscape and find flatter minima in adversarial training, which is generally regarded as an indicator of good generalization ability. After the flatness analysis, we show *how to optimize with randomized weights during adversarial training* in Sec. 4. A

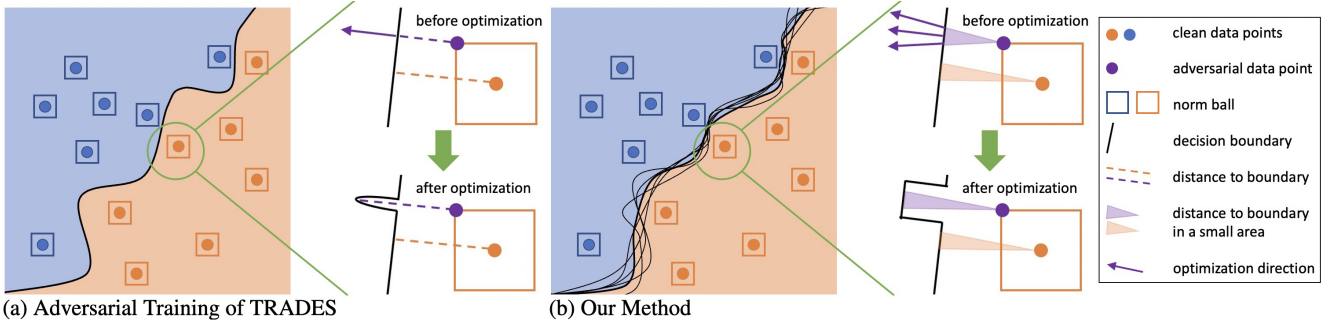


Figure 1. A conceptual illustration of decision boundaries learned via (a) adversarial training of TRADES and (b) our method. (a) shows that TRADES considers a deterministic model and optimizes the distance of adversarial data and boundary only through one direction. Our method in (b) takes into account randomized weights (perturbed boundaries) and optimizes the distance of adversarial data and boundary via multi directions in a small area. The boundary learned by our method can be smoother and more robust against new data.

novel adversarial training method based on TRADES [80] is proposed to reconcile adversarial robustness with clean accuracy by closing the gap between clean latent space and adversarial latent space over randomized weights. Specifically, we utilize Taylor series to expand the objective function over weights, in such a way that we can deconstruct the function into Taylor series (e.g., zeroth term, first term, second term, etc). From an algorithmic viewpoint, these Taylor terms can thus replace the objective function effectively and time-efficiently. As Fig. 1 shows, since our method takes randomized models into consideration during training, the learned boundary is smoother and the learned model is more robust in a small perturbed area.

We validate the effectiveness of our optimization method with the first and the second derivative terms of Taylor series. In consideration of training complexity and efficiency, we omit the third and higher derivative terms. Through an extensive set of experiments on a wide range of datasets (CIFAR-10 [40], CIFAR-100, SVHN [51]) and model architectures (ResNet [27], WideResNet [78], VGG [65], MobileNetV2 [62]), we find that *our method can further enhance the state-of-the-art adversarial training methods on both adversarial robustness and clean accuracy, consistently across the datasets and the model architectures.* Overall, this paper makes the following contributions:

- We conduct a pilot study of the trade-off between adversarial robustness and clean accuracy with randomized weights, and offer a new insight on the smoothed weights update and the flat minima during adversarial training (Sec. 3).
- We propose a novel adversarial training method under a randomized model to smooth the weights. The key enabler is the Taylor series expansion (in Sec. 4) of the robustness loss function over randomized weights (deterministic weights with random noise), so that the optimization can be simultaneously done over the zeroth, first, and second orders of Taylor series. In doing so, the proposed method can effectively enhance adversarial robustness without a significant compromise

on clean accuracy.

- An extensive set of empirical results are provided to demonstrate that our method can improve both robustness and clean accuracy consistently across different datasets and different network architectures (Sec. 5).

## 2. Preliminaries

**Basic Notation.** Consider the classification task with the training set  $\mathcal{S} = \{s_1, \dots, s_m\}$  where  $m$  samples are drawn from the data distribution  $\mathcal{D}$ . For notational convenience, we omit the label  $y$  of sample  $s$ . The adversarial example  $s'$  is generally not in the natural dataset  $\mathcal{S}$ , such that  $\|s' - s\|_p \leq \epsilon$  where  $\|\cdot\|_p$  is by default the  $\ell_p$ -norm. Let  $f_w(\cdot)$  be a learning model parameterized by  $w$ , where  $w$  is the vectorization of weight matrix  $W$ .  $\mathcal{L}(f_w(s), y)$  is the cross-entropy loss between  $f_w(s)$  and  $y$  with normalization. Define  $\mathcal{L}(f_w(\mathcal{S}), \mathcal{Y}) := \mathbb{E}_{s \in \mathcal{S}} [\mathcal{L}(f_w(s), y)]$ .

**Vanilla Adversarial Training.** Adversarial training can be formulated as a minimax optimization problem [45]

$$\min_w \left\{ \mathbb{E}_{s \sim \mathcal{D}} \left[ \max_{s': \|s' - s\|_p \leq \epsilon} \mathcal{L}(f_w(s'), y) \right] \right\}, \quad (1)$$

where  $s'$  is an adversarial example causing the largest loss within an  $\epsilon$ -ball centered at a clean example  $s$  with respect to a norm distance. Adversarial training aims to learn a model with the loss of adversarial examples minimized, without caring about the loss of the clean examples.

**TRADES-like Adversarial Training.** Advanced adversarial training methods, such as TRADES [80], strike a trade-off between clean accuracy and adversarial robustness, by optimizing the following problem

$$\min_w \left\{ \mathbb{E}_{s \sim \mathcal{D}} \left[ \mathcal{L}(f_w(s), y) + \max_{s': \|s' - s\|_p \leq \epsilon} \mathcal{L}(f_w(s), f_w(s')) / \lambda \right] \right\},$$

where the first term contributes to clean accuracy, and the second term with hyperparameter  $\lambda$  can be seen as a regularization for adversarial robustness that balances the outputs of clean and adversarial examples.

Nevertheless, each adversarial example  $s'$  is generated from a clean example  $s$  based on a specific weight  $w$ , and

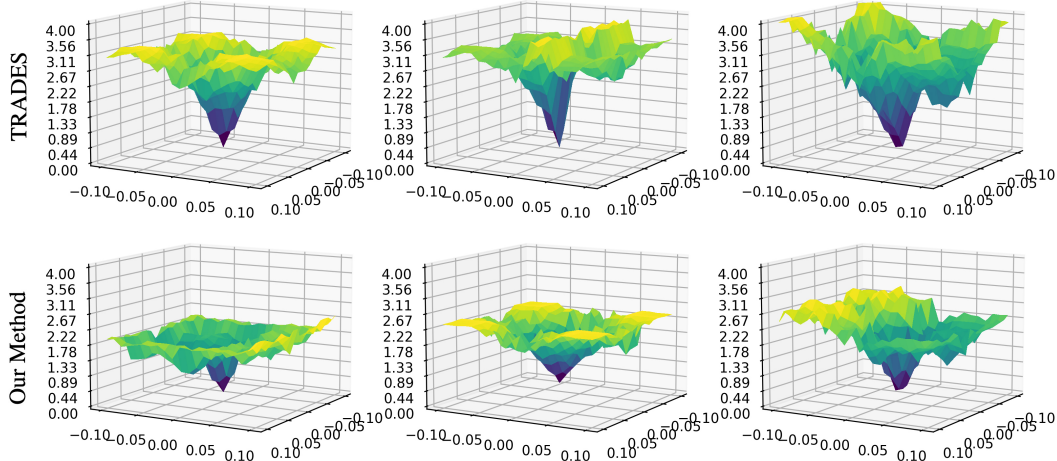


Figure 2. Comparison of loss landscapes of TRADES trained model (the first row) and TRADES+ $1_{st}+2_{nd}$  (our method, the second row) trained model. Loss plots in each column are generated from the same original image randomly chosen from the CIFAR-10 test dataset. Following the settings in [20], the  $z$  axis represents the loss, the  $x$  and  $y$  axes represent the magnitude of the perturbation added in the directions of  $\text{sign}\nabla_{\mathbf{s}}f(\mathbf{s})$  and Rademacher(0.5) respectively. We provide more empirical results in Appendix A.

the weight  $\mathbf{w}$  is obtained by optimizing both clean and adversarial data with SGD. This is a chick-and-egg problem. A not-very-well-trained model  $\mathbf{w}$  may generate some misleading adversarial data samples, which in turn drive the training process away from the optimum.

Inspired by the idea of model smoothing, we believe that the introduction of noise term into the training process may smooth the update of weights and thus robustify the models. To elaborate on this, we introduce our method from the perspective of flatness of loss landscape in Sec. 3. This is followed by a formalization in Sec. 4, where a noise term is injected to network weight  $\mathbf{w}$  during adversarial training to transform deterministic weights into randomized weights.

### 3. Motivation

Flatness is commonly thought to be a metric of standard generalization: the loss surface at the final learnt weights for well-generalizing models is generally “flat” [18, 21, 35, 52]. Moreover, [73] believed that a flatter adversarial loss landscape reduces the generalization gap in robustness. This aligns with [28], where the authors demonstrated that the local Lipschitz constant can be used to explicitly quantify the robustness of machine learning models. Many empirical defensive approaches, such as hessian/curvature-based regularization [48], gradient magnitude penalty [71], smoothening with random noise [44], or entropy regularization [33], have echoed the flatness performance of a robust model. However, all the above approaches require significant computational or memory resources, and many of them, such as hessian/curvature-based solutions, may suffer from standard accuracy decreases [26].

Stochastic weight averaging (SWA) [32] is known to find solutions that are far flatter than SGD, is incredibly simple to implement, enhances standard generalization, and

has almost no computational overhead. It was developed to guarantee weight smoothness by simply averaging various checkpoints throughout the training trajectory. SWA has been applied to semi-supervised learning, Bayesian inference, and low-precision training. Especially, [8] used SWA for the first time in adversarial training in order to smooth the weights and locate flatter minima, and showed that it enhances the adversarial generalization.

Different from SWA, we propose a novel method (with details provided in Sec. 4) via Taylor expansion of a small Gaussian noise to smooth the update of weights. This method regards the TRADES [80], a state-of-the-art adversarial training method, as a special case with only zero-th term, and is able to locate flatter minima. To demonstrate the latter, we visualize the loss landscape with respect to both input and weight spaces. As shown in Fig. 2, compared to the TRADES, our method significantly flattens the rugged landscape. In addition, it is worth noting that our empirical results (in Sec. 5) demonstrate that our method can further improve not only clean accuracy but also adversarial accuracy compared with other SOTA (includes SWA from [8]).

### 4. Methodology

Given a deterministic model of weight  $\mathbf{w}$ , we consider an additive small Gaussian noise  $\mathbf{u}$  to smooth  $\mathbf{w}$ . Based on TRADES [80], we propose that the robust optimization problem can be further enhanced by

$$\min_{\mathbf{w}} \mathbb{E}_{\mathbf{u}} \left[ \mathbb{E}_{\mathbf{s}} (\mathcal{L}(f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}), y)) + \mathbb{E}_{\mathbf{s}} \max_{\mathbf{s}': \|\mathbf{s}'-\mathbf{s}\|_p \leq \epsilon} \text{KL}(f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}) || f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}')) / \lambda \right], \quad (2)$$

where the first term contributes to clean accuracy, and the second term with hyperparameter  $\lambda$  can be seen as a regu-

---

**Algorithm 1** Adversarial training with randomized weights

---

**Input:** minibatch  $\{\mathbf{s}_i\}_{i=1}^n$ , network architecture parametrized by  $\mathbf{w}$ , learning rate  $\eta_l$ , step size  $\eta_s$ , hyper-parameters  $\eta, \lambda$ , zero mean Gaussian  $\mathbf{u}$ , number of iterations  $K$  in inner optimization.

**Output:** Robust network  $f_{\mathbf{w}}$

Randomly initialize network  $f_{\mathbf{w}}$ , or initialize network with pre-trained configuration

**repeat**

► Generate adversarial examples:

Sample  $u$  from  $\mathbf{u}$

**for**  $i = 1$  to  $n$  **do**

$\mathbf{s}'_i \leftarrow \mathbf{s}_i + 0.001 \cdot \mathcal{N}(\mathbf{0}, \mathbf{I})$

**for**  $k = 1$  to  $K$  **do**

$\mathbf{s}'_i \leftarrow \Pi(\eta_s \mathbf{sign}(\nabla_{\mathbf{s}'_i} \mathcal{L}(f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}_i), f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}'_i))) + \mathbf{s}'_i)$ , subject to  $\|\mathbf{s}_i - \mathbf{s}'_i\|_p \leq \epsilon$ , where  $\Pi$  is the projection operator.

**end for**

**end for**

► Optimization:  $\mathbf{w} \leftarrow \mathbf{w} - \eta_l \frac{1}{n} \sum_{i=1}^n \nabla_{\mathbf{w}} \left[ \mathcal{L}(f_{\mathbf{w}}(\mathbf{s}_i), y_i) \right. \\ \text{(zeroth term)} \quad \left. + \mathcal{L}(g_{\mathbf{s}_i}(\mathbf{w}), g_{\mathbf{s}'_i}(\mathbf{w})) / \lambda \right. \\ \text{(first term)} \quad \left. + \eta \mathbb{E}_{\mathbf{u}} (\mathcal{L}(g'_{\mathbf{s}_i}(\mathbf{w})^T \mathbf{u}, g'_{\mathbf{s}'_i}(\mathbf{w})^T \mathbf{u})) \right. \\ \left. \text{(second term)} \quad \left. + \frac{\eta}{2} \mathbb{E}_{\mathbf{u}} (\mathcal{L}(\mathbf{u}^T g''_{\mathbf{s}_i}(\mathbf{w}) \mathbf{u}, \mathbf{u}^T g''_{\mathbf{s}'_i}(\mathbf{w}) \mathbf{u})) \right]$ ,

where various optimization methods can be applied: **zeroth** term optimization (TRADES), **zeroth** + **first** terms optimization and **zeroth** + **first** + **second** terms optimization.

**until** training converged

---

larization for adversarial robustness with weights  $\mathbf{w}$  and a small Gaussian noise  $\mathbf{u}$ .

Following [17, 35], we let the Gaussian  $\mathbf{u}$  be small enough to maintain the generalization performance on clean data. Then, we can extend Eq. (2) by replacing  $\text{KL}(\cdot|\cdot)$  with a multi-class calibrated loss  $\mathcal{L}(\cdot, \cdot)$  [80]:

$$\begin{aligned} \min_{\mathbf{w}} \mathbb{E}_{\mathbf{u}} \left[ \mathbb{E}_{\mathbf{s}} (\mathcal{L}(f_{\mathbf{w}}(\mathbf{s}), y)) \right. \\ \left. + \mathbb{E}_{\mathbf{s}} \max_{\mathbf{s}': \|\mathbf{s}' - \mathbf{s}\|_p \leq \epsilon} \mathcal{L}(f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}), f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}')) / \lambda \right]. \end{aligned} \quad (3)$$

As the minimax optimization is entangled with expectation of randomized weights, it is challenging to solve such optimization problem directly. Instead, we propose to solve this problem in an alternating manner between the inner maximization and outer minimization.

For the inner maximization, with a given model weight  $\mathbf{w}$ , we solve it by adopting the commonly used gradient-based approach (e.g., PGD) to generate the adversarial perturbation. That is, at the  $t$ -th iteration, the adversarial perturbation is produced iteratively by letting

$$\mathbf{s}'_{t+1} \leftarrow \Pi \left( \mathbf{s}'_t + \eta_s \mathbf{sign}(\nabla_{\mathbf{s}'_t} \mathcal{L}(f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}), f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}'_t))) \right) \quad (4)$$

until the maximum allowed iterations are reached, where  $\Pi$  is the projection operator,  $\eta_s$  is the step size and  $u$  is a sample of  $\mathbf{u}$ . To reduce complexity, we only use one sample  $u$  to generate adversarial example for each minibatch. Nevertheless, sufficient samples of  $\mathbf{u}$  can still be produced in one

epoch as there are many minibatches. As Fig. 1 shows, the adversarial example is generated by one model (boundary), then the optimization with randomized weights in multiple directions can be executed through the following minimization method.

For the outer minimization, when the optimal perturbed sample  $\mathbf{s}'$  is generated by the inner maximization, we solve the following problem

$$\min_{\mathbf{w}} \mathbb{E}_{\mathbf{s}} \left[ \mathcal{L}(f_{\mathbf{w}}(\mathbf{s}), y) + \mathbb{E}_{\mathbf{u}} \mathcal{L}(f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}), f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}')) / \lambda \right]. \quad (5)$$

This optimization problem is challenging due to the entanglement between the deterministic weight  $\mathbf{w}$  and the random perturbation  $\mathbf{u}$ , which makes gradient updating much involved.

To resolve this issue, we decompose the two components in such a way that gradient updating and model averaging can be done separately. Let  $f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}) = g_{\mathbf{s}}(\mathbf{w} + \mathbf{u})$ , by Taylor expansion of  $g_{\mathbf{s}}(\mathbf{w} + \mathbf{u})$  at  $\mathbf{w}$ , we have

$$\begin{aligned} g_{\mathbf{s}}(\mathbf{w} + \mathbf{u}) &= g_{\mathbf{s}}(\mathbf{w}) + g'_{\mathbf{s}}(\mathbf{w})^T \mathbf{u} + \mathbf{u}^T \frac{g''_{\mathbf{s}}(\mathbf{w})}{2!} \mathbf{u} \\ &\quad + \mathcal{O}_{\mathbf{s}}(\|\mathbf{u}\|^2), \end{aligned} \quad (6)$$

where  $g'_{\mathbf{s}}(\mathbf{w}) = \frac{\partial g_{\mathbf{s}}(\mathbf{w})}{\partial \mathbf{w}} \in \mathbb{R}^{d \times 1}$  and  $g''_{\mathbf{s}}(\mathbf{w}) = \frac{\partial^2 g_{\mathbf{s}}(\mathbf{w})}{\partial \mathbf{w} \partial \mathbf{w}} \in \mathbb{R}^{d \times d}$  are the first and second derivative of the function  $g_{\mathbf{s}}(\mathbf{w})$  versus the weight  $\mathbf{w}$ , respectively, and  $\mathcal{O}_{\mathbf{s}}(\|\mathbf{u}\|^2)$  tends to be zero with respect to  $\|\mathbf{u}\|^2$ . By such approximation, the gradient computation will be solely done on the

deterministic part  $\mathbf{w}$ , and model smoothing with averaged random variable  $\mathbf{u}$  is done independently.

The intuition behind is the following. For the randomized model  $\mathbf{w} + \mathbf{u}$  with a specific input  $\mathbf{s}$ , the Taylor approximation explicitly takes into account the deterministic model  $g_{\mathbf{s}}(\mathbf{w})$ , the projection of its gradient onto the random direction  $g'_{\mathbf{s}}(\mathbf{w})^T \mathbf{u}$ , and the projection of its curvature onto the subspace spanned by  $\mathbf{u}^T \frac{g''_{\mathbf{s}}(\mathbf{w})}{2!} \mathbf{u}$ , capturing higher-order information of the loss function with respect to  $\mathbf{u}$ .

To further simplify the computation, we minimize an upper bound of Eq. (5) instead. Due to the local convexity of loss landscape, by Jensen’s inequality, we conclude that

$$\mathcal{L}\left(\sum_i x_i, \sum_i y_i\right) \leq \sum_i \mathcal{L}(x_i, y_i). \quad (7)$$

Therefore, the second term in Eq. (5) can be upper bounded by

$$\begin{aligned} & \mathbb{E}_{\mathbf{u}} \mathcal{L}(f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}), f_{\mathbf{w}+\mathbf{u}}(\mathbf{s}')) \\ & \leq \mathcal{L}(g_{\mathbf{s}}(\mathbf{w}), g_{\mathbf{s}'}(\mathbf{w})) + \mathbb{E}_{\mathbf{u}} (\mathcal{L}(g'_{\mathbf{s}}(\mathbf{w})^T \mathbf{u}, g'_{\mathbf{s}'}(\mathbf{w})^T \mathbf{u})) \\ & \quad + \frac{1}{2} \mathbb{E}_{\mathbf{u}} (\mathcal{L}(\mathbf{u}^T g''_{\mathbf{s}}(\mathbf{w}) \mathbf{u}, \mathbf{u}^T g''_{\mathbf{s}'}(\mathbf{w}) \mathbf{u})), \end{aligned} \quad (8)$$

where the terms on the right-hand side are referred to as the zeroth, first, and second order Taylor expressions. Instead of minimizing Eq. (5) directly, we minimize the upper bound in Eq. (8), which is easier to compute in practical models. Details of the optimization are given in Appendix D.

The pseudo-code of the proposed adversarial training method is presented in Algorithm 1. Note that the zeroth term optimization in Algorithm 1 is almost the same as TRADES [80], thus we use TRADES and AWP-TRADES [74] as (zeroth term optimization) baselines in our experiments in Sec. 5, and verify how much first and second terms optimization can improve. Note that we perform both AWP and our method through applying randomized weight perturbations on AWP updated weights.

For all of our experiments, we limit the noise variance of  $\mathbf{u}$  to be small enough [17, 35] to maintain the training, e.g.,  $\sigma = 0.01$ . For the zeroth term hyper-parameter  $1/\lambda$  in Algorithm 1, we set  $1/\lambda = 6$  for all of our experiments, which is a widely-used setting for TRADES [53, 74, 80]. The number of inner optimization iterations  $K$  is set to 10 as usual. To make our method more flexible, we set  $\eta$  and  $\frac{\eta}{2}$  as the first term and second term hyper-parameters, respectively, and analyze the sensitivity of  $\eta$  in Sec. 5.1.

The main novelty of our proposed method is two-fold: **(1)** Instead of considering a single model with a deterministic weight  $\mathbf{w}$ , we consider a model ensemble with  $\mathbf{w} + \mathbf{u}$  to smooth the update of weights. By averaging these models during adversarial training, we come up with a robust model with smoothed classification boundary tolerant to potential adversarial weight perturbation. **(2)** When averaging

over the ensemble of randomized models during training, we disentangle gradient from random weight perturbation so that the gradient updating can be computed efficiently. In particular, we apply Taylor expansion at the mean weight  $\mathbf{w}$  (i.e., the deterministic component) and approximate the cross-entropy loss function with the zeroth, first, and second Taylor terms. In doing so, the deterministic and statistical components of  $\mathbf{w} + \mathbf{u}$  can be decomposed and then computed independently.

## 5. Empirical results

In this section, we first discuss the hyper-parameter sensitivity of our method, and then evaluate the robustness on benchmark datasets against various white-box, black-box attacks and Auto Attack with  $\ell_2$  and  $\ell_\infty$  threat models.

**Adversarial Training Setting.** We train PreAct ResNet-18 [27] for  $\ell_\infty$  and  $\ell_2$  threat models on CIFAR-10/100 [40] and SVHN [51]. In addition, we also train WideResNet-34-10 [78] for CIFAR-10/100, VGG16 and MobileNetV2 for CIFAR-10, with  $\ell_\infty$  threat model. We adopt the widely used adversarial training setting [60]: for the  $\ell_\infty$  threat model,  $\epsilon = 8/255$  and step size  $2/255$ ; for the  $\ell_2$  threat model,  $\epsilon = 128/255$  and step size  $15/255$ . For normal adversarial training, the training examples are generated with 10 steps. All models (except SVHN) are trained for 200 epochs using SGD with momentum 0.9, batchsize 128, weight decay  $5 \times 10^{-4}$ , and an initial learning rate of 0.1 that is divided by 10 at the 100th and 150th epochs. Except for setting the starting learning rate to 0.01 for SVHN, we utilize the same other settings. Simple data augmentations are used, such as  $32 \times 32$  random crop with 4-pixel padding and random horizontal flip. We report the highest robustness that ever achieved at different checkpoints for each dataset and report the clean accuracy on the model which gets the highest PGD-20 accuracy. We omit the standard deviations of 3 runs as they are very small ( $< 0.40\%$ ) and implement all models on NVIDIA A100.

**Evaluation Setting.** We evaluate the robustness with white-box attacks, black-box attacks and auto attack. For **white-box attacks**, we adopt PGD-20 [45] and CW-20 [4] (the  $\ell_\infty$  version of CW loss optimized by PGD-20) to evaluate trained models. For **black-box attacks**, we generate adversarial perturbations by attacking a surrogate normal adversarial training model (with same setting) [55], and then apply these adversarial examples to the defense model and evaluate the performances. The attacking methods for black-box attacks we have used are PGD-20 and CW-20. For **Auto Attack (AA)** [10], one of the strongest attack methods, we adopt it through a mixture of different parameter-free attacks which include three white-box attacks (APGD-CE [10], APGD-DLR [10], and FAB [9]) and one black-box attack (Square Attack [2]). We provide more details about the experimental setups in Appendix B.

Table 1. First and Second Derivative terms optimization on CIFAR-10 with  $\ell_\infty$  threat model for PreAct ResNet18. Classification accuracy (%) on clean images and under PGD-20 attack, CW-20 attack and Auto Attack with different hyper-parameters  $\eta = 0.05, 0.1, 0.2, 0.3, 0.4$ . We highlight the best results in **bold**.

Method	Clean	PGD-20	CW-20	AA	Method	Clean	PGD-20	CW-20	AA
$1_{st}$ (0.05)	83.19	53.98	52.12	48.2	$1_{st}+2_{nd}$ (0.05)	83.25	54.07	51.93	48.4
$1_{st}$ (0.1)	82.86	54.29	52.24	49.3	$1_{st}+2_{nd}$ (0.1)	83.47	54.42	52.28	49.7
$1_{st}$ (0.2)	83.55	54.86	<b>52.65</b>	48.8	$1_{st}+2_{nd}$ (0.2)	84.13	<b>54.91</b>	<b>52.53</b>	<b>50.3</b>
$1_{st}$ (0.3)	<b>83.96</b>	<b>55.05</b>	52.54	<b>49.7</b>	$1_{st}+2_{nd}$ (0.3)	<b>84.27</b>	54.36	51.80	48.9
$1_{st}$ (0.4)	83.93	54.65	52.23	48.6	$1_{st}+2_{nd}$ (0.4)	84.14	54.38	51.56	49.6

Table 2. First and Second Derivative terms optimization on CIFAR-10/CIFAR-100 with  $\ell_\infty$  threat model for WideResNet, compared with current state-of-the-art. Classification accuracy (%) on clean images and under PGD-20 attack, CW-20 attack ( $\epsilon = 0.031$ ) and Auto Attack ( $\epsilon = 8/255$ ). The results of our methods are in **bold**. Note that \* is under PGD-40 attack and \*\* is under PGD-10 attack.

Dataset	Method	Architecture	Clean	PGD-20	CW-20	AA
CIFAR-10 $\ell_\infty$	Lee et al. (2020) [42]	WRN-34-10	92.56	59.75	54.53	39.70
	Wang et al. (2020) [72]	WRN-34-10	83.51	58.31	54.33	51.10
	Rice et al. (2020) [60]	WRN-34-20	85.34	-	-	53.42
	Zhang et al. (2020) [81]	WRN-34-10	84.52	-	-	53.51
	Pang et al. (2021) [54]	WRN-34-20	86.43	57.91**	-	54.39
	Jin et al. (2022) [36]	WRN-34-20	86.01	61.12	57.93	55.90
	Gowal et al. (2020) [24]	WRN-70-16	85.29	58.22*	-	57.20
	Zhang et al. (2019) [80] (0 <sub>th</sub> )	WRN-34-10	84.65	56.68	54.49	53.0
	+ Ours ( $1_{st}$ )	WRN-34-10	<b>85.51</b>	<b>58.34</b>	<b>56.06</b>	<b>54.0</b>
	+ Ours ( $1_{st}+2_{nd}$ )	WRN-34-10	<b>85.98</b>	<b>58.47</b>	<b>56.13</b>	<b>54.2</b>
CIFAR-100 $\ell_\infty$	Wu et al. (2020) [74] (0 <sub>th</sub> )	WRN-34-10	85.17	59.64	57.33	56.2
	+ Ours ( $1_{st}$ )	WRN-34-10	<b>86.10</b>	<b>61.47</b>	<b>58.09</b>	<b>57.1</b>
	+ Ours ( $1_{st}+2_{nd}$ )	WRN-34-10	<b>86.12</b>	<b>61.45</b>	<b>58.22</b>	<b>57.4</b>
	Cui et al. (2021) [11]	WRN-34-10	60.43	35.50	31.50	29.34
	Gowal et al. (2020) [24]	WRN-70-16	60.86	31.47*	-	30.03
	Zhang et al. (2019) [80] (0 <sub>th</sub> )	WRN-34-10	60.22	32.11	28.93	26.9
	+ Ours ( $1_{st}$ )	WRN-34-10	<b>63.01</b>	<b>33.26</b>	<b>29.44</b>	<b>28.1</b>
	+ Ours ( $1_{st}+2_{nd}$ )	WRN-34-10	<b>62.93</b>	<b>33.36</b>	<b>29.61</b>	<b>27.9</b>
	Wu et al. (2020) [74] (0 <sub>th</sub> )	WRN-34-10	60.38	34.09	30.78	28.6
	+ Ours ( $1_{st}$ )	WRN-34-10	<b>63.98</b>	<b>35.36</b>	<b>31.63</b>	<b>29.8</b>
+ Ours ( $1_{st}+2_{nd}$ )	WRN-34-10	<b>64.71</b>	<b>35.73</b>	<b>31.41</b>	<b>30.2</b>	

## 5.1. Sensitivity of hyper-parameter

In our proposed algorithm, the regularization hyper-parameter  $\eta$  is crucial. We use numerical experiments on CIFAR-10 with PreAct ResNet-18 to illustrate how the regularization hyper-parameter influences the performance of our robust classifiers. To develop robust classifiers for multi-class tasks, we use the gradient descent update formula in Algorithm 1, with  $\mathcal{L}$  as the cross-entropy loss. Note that, for the zeroth term hyper-parameter  $1/\lambda$  in Algorithm 1, we set  $1/\lambda = 6$  for all of our experiments, which is a widely used setting for TRADES. We train the models with  $\eta = 0.05, 0.1, 0.2, 0.3, 0.4$  respectively. All models are trained with 200 epochs and 128 batchsize.

In Tab. 1, we observe that as the regularization parameter  $\eta$  increases, the clean accuracy and robust accuracy almost both go up first and then down. In particular, the accuracy under Auto Attack is sensitive to the regularization

hyper-parameter  $\eta$ . Considering both robustness accuracy and clean accuracy, it is not difficult to find that  $\eta = 0.3$  is the best for first term optimization and  $\eta = 0.2$  is the best for first + second terms optimization. Thus in the following experiments, we set  $\eta = 0.3$  for first term optimization and  $\eta = 0.2$  for first + second terms optimization as default.

## 5.2. Comparison with SOTA on WideResNet

In Tab. 2, we compare our method with state-of-the-art on WideResNet with CIFAR-10 and CIFAR-100. For zeroth term optimization, we use TRADES [80] and AWP-TRADES [74] as baselines with  $1/\lambda = 6$ . We also report other state-of-the-art, include AVMixup [42], MART [72], [60], [81], [54], TRADES+LGGAT [11] and [24].

The results in Tab. 2 demonstrate that our method can both improve clean accuracy and robustness accuracy consistently over different datasets and models. Especially for

Table 3. Adversarial training across datasets on PreAct ResNet18 with  $\ell_2$  threat model. Classification accuracy (%) on clean images and under PGD-20 attack and Auto Attack. The results of our methods are in **bold**.

Method	CIFAR-10			CIFAR-100			SVHN		
	Clean	PGD	AA	Clean	PGD	AA	Clean	PGD	AA
Wu et al. (2020) [74](0 <sub>th</sub> )	87.05	72.08	71.6	62.87	45.11	41.8	92.86	72.45	63.8
+ Ours ( $1_{st}+2_{nd}$ )	<b>88.41</b>	<b>72.85</b>	<b>72.0</b>	<b>65.59</b>	<b>46.88</b>	<b>42.4</b>	<b>93.98</b>	<b>73.07</b>	<b>64.5</b>

Table 4. Adversarial training across datasets on PreAct ResNet18 with  $\ell_\infty$  threat model. Classification accuracy (%) on clean images and black-box attacks. Black-box adversarial examples are generated by a surrogate normal adversarial training model (of same setting) with PGD-20 attack and CW-20 attack. The results of our methods are in **bold**.

Method	CIFAR-10			CIFAR-100			SVHN		
	Clean	PGD	CW	Clean	PGD	CW	Clean	PGD	CW
Wu et al. (2020) [74](0 <sub>th</sub> )	82.78	61.79	59.42	58.33	38.55	36.70	93.77	63.80	59.54
+ Ours ( $1_{st}+2_{nd}$ )	<b>84.27</b>	<b>62.56</b>	<b>60.58</b>	<b>61.08</b>	<b>39.82</b>	<b>37.84</b>	<b>95.11</b>	<b>66.16</b>	<b>63.59</b>

Table 5. Adversarial training across VGG16, MobileNetV2 on CIFAR-10 with  $\ell_\infty$  threat model. Classification accuracy (%) on clean images and under PGD-20 attack, CW-20 attack and Auto Attack. The results of our methods are in **bold**.

Method	VGG16				MobileNetV2			
	Clean	PGD-20	CW-20	AA	Clean	PGD-20	CW-20	AA
Zhang et al. (2019) [80] (0 <sub>th</sub> )	79.78	49.88	46.95	44.3	79.73	51.41	48.43	46.4
+ Ours ( $1_{st}+2_{nd}$ )	<b>80.99</b>	<b>50.13</b>	<b>47.09</b>	<b>44.4</b>	<b>81.86</b>	<b>53.34</b>	<b>49.93</b>	<b>47.8</b>
Wu et al. (2020) [74] (0 <sub>th</sub> )	78.46	51.19	47.41	46.3	79.86	53.56	50.11	47.7
+ Ours ( $1_{st}+2_{nd}$ )	<b>80.31</b>	<b>52.71</b>	<b>48.38</b>	<b>46.5</b>	<b>81.95</b>	<b>55.37</b>	<b>51.55</b>	<b>49.4</b>

Table 6. Time consumption and GPU memory usage for WideResNet on CIFAR-10 with  $\ell_\infty$  threat model. We deploy each model on a single NVIDIA A100 with batchsize 128. The results of our methods are in **bold**.

Method	WideResNet-34-10	
	Time/Epoch	GPU Memory
Zhang et al. (2019) [80] (0 <sub>th</sub> )	1666s	12875MB
+ Ours ( $1_{st}$ )	<b>2161s</b>	<b>20629MB</b>
+ Ours ( $1_{st}+2_{nd}$ )	<b>2362s</b>	<b>26409MB</b>

Auto Attack, our method can get 57.4% on CIFAR-10 and 30.2% on CIFAR-100 with WRN-34-10, even surpass the performance of WRN-70-16 from [24].

### 5.3. Other empirical results

**Adversarial training with  $\ell_2$  threat model.** For the experiments with  $\ell_2$  threat model on CIFAR-10, CIFAR-100 and SVHN in Tab. 3, the results still support that our method can enhance the performance under clean data, PGD-20 attack and Auto Attack. For CIFAR-100, it can even increase about 3% on clean accuracy.

**Robustness under black-box attacks.** We train PreAct ResNet18 for  $\ell_\infty$  threat model on CIFAR-10, CIFAR-100 and SVHN. The black-box adversarial examples are generated by a surrogate normal adversarial training model (of same setting) with PGD-20 attack and CW-20 attack. Tab. 4 shows our method is also effective under black-box attacks. For SVHN, ours can even obtain a significant improvement over the existing ones.

**Robustness on other architectures.** We train VGG16

and MobileNetV2 for  $\ell_\infty$  threat model on CIFAR-10. Our results in Tab. 5 show a comprehensive improvement under clean data and robustness accuracy. Particularly, for MobileNetV2, ours can achieve an improvement greater than 1% and 2% under Auto Attack and clean data, respectively.

**Comparison with SWA from [8].** The best PGD-20 accuracy and AA accuracy on CIFAR-10 for ResNet-18 from [8] are 52.14% and 49.44% respectively, whereas ours are 55.13% and 50.8% in Tab. 1 (with  $1_{st}+2_{nd}$  (0.2)).

More empirical results are given in Appendix E.

### 5.4. Limitations and future work

We use some approximation methods, e.g., Eq. (8), to reduce the complexity of our adversarial training method. The results in Tab. 6 show that though our method with WideResNet-34-10 can work on a single NVIDIA A100, the growth of training time and GPU memory still cannot be ignored. In the future work, we plan to further reduce the complexity of our algorithm and apply the first and second terms optimization on larger datasets.

## 6. Related work

### 6.1. Adversarial training

Adversarial training methods can generally be divided into three groups. In the first group, Eq. (1) is translated into an equivalent, (or approximate, expression), which is mainly used to narrow the distance between  $f_w(s)$  and  $f_w(s')$ . For example, ALP [19, 38] estimates the similar-

ity between  $f_w(s)$  and  $f_w(s')$ , and maximizes this similarity in the objective function. MMA [13] proposes each correctly classified instance  $s$  to leave sufficiently the decision boundary, through making the size of shortest successful perturbation as large as possible. TRADES [80] looks for the adversarial example  $s'$  which can obtain the largest KL divergence between  $f_w(s)$  and  $f_w(s')$ , then trains with these adversarial examples. Besides, [84] adopts the similarity between local distributions of natural and adversarial example, [67] measures the similarity of local distributions with Wasserstein distance, and [15, 16, 46, 54] try to optimize with the distribution over a series of adversarial examples for a single input.

In the second group, adversarial examples are pre-processed before being used for training rather than being directly generated by attack algorithms. For example, label smoothing [8, 69] replaces the hard label  $y$  with the soft label  $\tilde{y}$ , where  $\tilde{y}$  is a combination of the hard label  $y$  and uniform distribution. A further empirical exploration with of how label smoothing works is provided in [49]. AVMixup [42, 83] defines a virtual sample in the adversarial direction based on these. Through linear interpolation of the virtual sample and the clean sample, it extends the training distribution with soft labels. Instead of smoothing labels, [79] perturbs the local neighborhoods with an unsupervised way to produce adversarial examples. In addition, data augmentation [25, 59] is also shown to be an invaluable aid for adversarial training.

The above two groups merely adjust the component parts of the min-max formalism. With AWP [74], one more maximization is performed to find a weight perturbation based on the generated adversarial examples. The altered weights [12] are then used in the outer minimization function to reduce the loss caused by the adversarial cases.

Different from previous work, this paper optimizes the distance between  $f_{w+u}(s)$  and  $f_{w+u}(s')$  through randomized weights, and therefore, effectively, it considers the optimization problem where the decision boundary is smoothed by these randomized models.

## 6.2. Randomized weights

In the previous work to study generalization or robustness, modeling neural network weights as random variables has been frequently used. For example, [50, 75] estimate the mutual information between random weights and dataset, provide the expected generalization bound over weight distribution. Another example of random weights is on the PAC-Bayesian framework, a well-known theoretical tool to bound the generalization error of machine learning models [41, 47, 56, 64]. In the recent years, PAC-Bayes is also developed to bound the generalization error or robustness error of deep neural networks [18, 21, 35, 52]. In addition, [23] draws random noise into deterministic weights, to measure

the mutual information between input dataset and activations under information bottleneck theory.

While we also consider randomized weights, a major difference from previous theoretical works is that randomized weights are applied to robustness, from which we find an insight to design a novel adversarial training method which decomposes objective function (parameterized with randomized weights) with Taylor series.

## 6.3. Flatness

The generalization performance of machine learning models is believed to be correlated with the flatness of loss curve [30, 31], especially for deep neural networks [6, 14, 37, 43, 76]. For instance, in pioneering research [30, 31], the minimum description length (MDL) [61] was used to demonstrate that the flatness of the minimum is a reasonable generalization metric to consider. [39] investigated the cause of the decrease in generalization for large-batch networks and demonstrated that large-batch algorithms tend to converge to sharp minima, resulting in worse generalization. [35] undertook a large-scale empirical study and found that flatness-based measurements correlate with generalization more strongly than weight norms and (margin and optimization)-based measures. [22] introduced Sharpness-Aware Minimization (SAM) to locate parameters that lie in neighborhoods with consistently low loss. Moreover, [8] used SWA [32] in adversarial training to smooth the weights and find flatter minima, and showed that this can improve adversarial generalization.

Different from previous work, this paper uses Taylor expansion of the injected Gaussian to smooth the update of weights and search for flat minima during adversarial training. Our empirical results show the effectiveness of our method to flatten loss landscape and improve both robustness and clean accuracy.

## 7. Conclusion

This work studies the trade-off between robustness and clean accuracy through the lens of randomized weights. Through empirical analysis of loss landscape, algorithmic design (via Taylor expansion) on training optimization, and extensive experiments, we demonstrate that optimizing over randomized weights can consistently outperform the state-of-the-art adversarial training methods not only in clean accuracy but also in robustness.

**Acknowledgment.** GJ is supported by the CAS Project for Young Scientists in Basic Research, Grant No.YSBR-040. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 956123, and is also supported by the U.K. EPSRC through End-to-End Conceptual Guarding of Neural Architectures [EP/T026995/1].



## References

- [1] Jean-Baptiste Alayrac, Jonathan Uesato, Po-Sen Huang, Al-hussein Fawzi, Robert Stanforth, and Pushmeet Kohli. Are labels required for improving adversarial robustness? In *NeurIPS*, 2019. 1
- [2] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: a query-efficient black-box adversarial attack via random search. In *ECCV*, 2020. 5
- [3] Yogesh Balaji, Tom Goldstein, and Judy Hoffman. Instance adaptive adversarial training: Improved accuracy tradeoffs in neural nets. *CoRR*, 2019. 1
- [4] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017. 5
- [5] Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C. Duchi, and Percy Liang. Unlabeled data improves adversarial robustness. In *NeurIPS*, 2019. 1
- [6] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer T. Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. In *International Conference on Learning Representations*, 2017. 8
- [7] Jinghui Chen and Quanquan Gu. Rays: A ray searching method for hard-label adversarial attack. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1739–1747, 2020. 15
- [8] Tianlong Chen, Zhenyu Zhang, Sijia Liu, Shiyu Chang, and Zhangyang Wang. Robust overfitting may be mitigated by properly learned smoothing. In *ICLR*, 2020. 1, 3, 7, 8
- [9] Francesco Croce and Matthias Hein. Minimally distorted adversarial examples with a fast adaptive boundary attack. In *ICML*, 2020. 5
- [10] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *ICML*, 2020. 5
- [11] Jiequan Cui, Shu Liu, Liwei Wang, and Jiaya Jia. Learnable boundary guided adversarial training. In *ICCV*, 2021. 1, 6
- [12] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 8
- [13] Gavin Weiguang Ding, Yash Sharma, Kry Yik Chau Lui, and Ruitong Huang. Mma training: Direct input space margin maximization through adversarial training. In *ICLR*, 2020. 8
- [14] Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. In *ICML*, 2017. 8
- [15] Yinpeng Dong, Zhijie Deng, Tianyu Pang, Jun Zhu, and Hang Su. Adversarial distributional training for robust deep learning. In *NeurIPS*, 2020. 8
- [16] Yinpeng Dong, Qi-An Fu, Xiao Yang, Tianyu Pang, Hang Su, Zihao Xiao, and Jun Zhu. Benchmarking adversarial robustness on image classification. In *CVPR*, 2020. 8
- [17] Gintare Karolina Dziugaite, Alexandre Drouin, Brady Neal, Nitarshan Rajkumar, Ethan Caballero, Linbo Wang, Ioannis Mitliagkas, and Daniel M. Roy. In search of robust measures of generalization. In *NeurIPS*, 2020. 4, 5
- [18] Gintare Karolina Dziugaite and Daniel M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *UAI*, 2017. 1, 3, 8
- [19] Logan Engstrom, Andrew Ilyas, and Anish Athalye. Evaluating and understanding the robustness of adversarial logit pairing. *arXiv preprint arXiv:1807.10272*, 2018. 1, 7
- [20] Logan Engstrom, Andrew Ilyas, and Anish Athalye. Evaluating and understanding the robustness of adversarial logit pairing. *CoRR*, 2018. 3, 13
- [21] Farzan Farnia, Jesse M. Zhang, and David Tse. Generalizable adversarial training via spectral normalization. In *ICLR*, 2019. 3, 8
- [22] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2021. 8
- [23] Ziv Goldfeld, Ewout van den Berg, Kristjan H. Greenewald, Igor Melnyk, Nam Nguyen, Brian Kingsbury, and Yury Polyanskiy. Estimating information flow in deep neural networks. In *ICML*, 2019. 8
- [24] Sven Gowal, Chongli Qin, Jonathan Uesato, Timothy A. Mann, and Pushmeet Kohli. Uncovering the limits of adversarial training against norm-bounded adversarial examples. *CoRR*, 2020. 6, 7
- [25] Sven Gowal, Sylvestre-Alvise Rebuffi, Olivia Wiles, Florian Stimberg, Dan Andrei Calian, and Timothy A. Mann. Improving robustness using generated data. In *NeurIPS*, 2021. 8
- [26] Vipul Gupta, Santiago Aklé Serrano, and Dennis DeCoste. Stochastic weight averaging in parallel: Large-batch training that generalizes well. In *ICLR*, 2020. 3
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 5
- [28] Matthias Hein and Maksym Andriushchenko. Formal guarantees on the robustness of a classifier against adversarial manipulation. In *NIPS*, 2017. 3
- [29] Dan Hendrycks, Kimin Lee, and Mantas Mazeika. Using pre-training can improve model robustness and uncertainty. In *ICML*, 2019. 1
- [30] Sepp Hochreiter and Jürgen Schmidhuber. Simplifying neural nets by discovering flat minima. In Gerald Tesauro, David S. Touretzky, and Todd K. Leen, editors, *Advances in Neural Information Processing Systems*, pages 529–536. MIT Press, 1994. 8
- [31] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural Comput.*, 9(1):1–42, 1997. 8
- [32] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry P. Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *UAI*, 2018. 1, 3, 8
- [33] Gauri Jagatap, Ameja Joshi, Animesh Basak Chowdhury, Siddharth Garg, and Chinmay Hegde. Adversarially robust learning via entropic regularization. *Frontiers Artif. Intell.*, 4:780843, 2021. 3

- [34] Adel Javanmard, Mahdi Soltanolkotabi, and Hamed Hassani. Precise tradeoffs in adversarial training for linear regression. In *COLT*, 2020. 1
- [35] Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic generalization measures and where to find them. In *ICLR*, 2020. 3, 4, 5, 8
- [36] Gaojie Jin, Xinping Yi, Wei Huang, Sven Schewe, and Xiaowei Huang. Enhancing adversarial training with second-order statistics of weights. In *CVPR*, 2022. 1, 6
- [37] Gaojie Jin, Xinping Yi, Liang Zhang, Lijun Zhang, Sven Schewe, and Xiaowei Huang. How does weight correlation affect generalisation ability of deep neural networks? In *NeurIPS*, 2020. 8
- [38] Harini Kannan, Alexey Kurakin, and Ian Goodfellow. Adversarial logit pairing. *arXiv preprint arXiv:1803.06373*, 2018. 1, 7
- [39] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*, 2017. 8
- [40] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 2, 5
- [41] John Langford and John Shawe-Taylor. Pac-bayes & margins. In *NeurIPS*, 2002. 8
- [42] Saehyung Lee, Hyungyu Lee, and Sungroh Yoon. Adversarial vertex mixup: Toward better adversarially robust generalization. In *CVPR*, 2020. 1, 6, 8
- [43] Tengyuan Liang, Tomaso A. Poggio, Alexander Rakhlin, and James Stokes. Fisher-rao metric, geometry, and complexity of neural networks. In *AISTATS*, 2019. 8
- [44] Xuanqing Liu, Minhao Cheng, Huan Zhang, and Cho-Jui Hsieh. Towards robust neural networks via random self-ensemble. In *ECCV*, 2018. 3
- [45] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018. 2, 5
- [46] Chengzhi Mao, Ziyuan Zhong, Junfeng Yang, Carl Vondrick, and Baishakhi Ray. Metric learning for adversarial robustness. In *NeurIPS*, 2019. 8
- [47] David A. McAllester. Pac-bayesian model averaging. In *COLT*, 1999. 8
- [48] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Jonathan Uesato, and Pascal Frossard. Robustness via curvature regularization, and vice versa. In *CVPR*, 2019. 3
- [49] Rafael Müller, Simon Kornblith, and Geoffrey Hinton. When does label smoothing help? In *NeurIPS*, 2019. 8
- [50] Jeffrey Negrea, Mahdi Haghifam, Gintare Karolina Dziugaite, Ashish Khisti, and Daniel M Roy. Information-theoretic generalization bounds for sgld via data-dependent estimates. In *NeurIPS*, 2019. 8
- [51] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. 2, 5
- [52] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. In *NeurIPS*, 2017. 1, 3, 8
- [53] Tianyu Pang, Min Lin, Xiao Yang, Jun Zhu, and Shuicheng Yan. Robustness and accuracy could be reconcilable by (proper) definition. *CoRR*, 2022. 5
- [54] Tianyu Pang, Xiao Yang, Yinpeng Dong, Hang Su, and Jun Zhu. Bag of tricks for adversarial training. In *ICLR*, 2021. 6, 8
- [55] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *AsiaCCS*, 2017. 5
- [56] Emilio Parrado-Hernández, Amiran Ambroladze, John Shawe-Taylor, and Shiliang Sun. Pac-bayes bounds with data dependent priors. *J. Mach. Learn. Res.*, 2012. 8
- [57] Rahul Rade and Seyed-Mohsen Moosavi-Dezfooli. Reducing excessive margin to achieve a better accuracy vs. robustness trade-off. In *International Conference on Learning Representations*, 2022.
- [58] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John C. Duchi, and Percy Liang. Understanding and mitigating the tradeoff between robustness and accuracy. In *ICML*, 2020. 1
- [59] Sylvestre-Alvise Rebuffi, Sven Gowal, Dan A. Calian, Florian Stimberg, Olivia Wiles, and Timothy A. Mann. Fixing data augmentation to improve adversarial robustness. *CoRR*, 2021. 8
- [60] Leslie Rice, Eric Wong, and J. Zico Kolter. Overfitting in adversarially robust deep learning. In *ICML*, 2020. 1, 5, 6
- [61] Jorma Rissanen. Modeling by shortest data description. *Autom.*, 14(5):465–471, 1978. 8
- [62] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018. 2
- [63] Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, and Aleksander Madry. Adversarially robust generalization requires more data. In *NeurIPS*, 2018. 1
- [64] Matthias W. Seeger. Pac-bayesian generalisation error bounds for gaussian process classification. *J. Mach. Learn. Res.*, 2002. 8
- [65] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 2
- [66] Gaurang Sriramanan, Sravanti Addepalli, Arya Baburaj, et al. Guided adversarial attack for evaluating and enhancing adversarial defenses. *Advances in Neural Information Processing Systems*, 33:20297–20308, 2020. 15
- [67] Matthew Staib and Stefanie Jegelka. Distributionally robust deep learning as a generalization of adversarial training. *NIPS workshop on Machine Learning and Computer Security*, 2017. 8
- [68] Dong Su, Huan Zhang, Hongge Chen, Jinfeng Yi, Pin-Yu Chen, and Yupeng Gao. Is robustness the cost of accuracy? - A comprehensive study on the robustness of 18 deep image classification models. In *ECCV*, 2018. 1
- [69] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016. 8
- [70] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. In *ICLR*, 2019. 1

- [71] Jianyu Wang and Haichao Zhang. Bilateral adversarial training: Towards fast training of more robust models against adversarial attacks. In *ICCV*, 2019. 3
- [72] Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. Improving adversarial robustness requires revisiting misclassified examples. In *ICLR*, 2020. 6
- [73] Dongxian Wu, Yisen Wang, and Shutao Xia. Revisiting loss landscape for adversarial robustness. *CoRR*, 2020. 3
- [74] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. In *NeurIPS*, 2020. 1, 5, 6, 7, 8
- [75] Aolin Xu and Maxim Raginsky. Information-theoretic analysis of generalization capability of learning algorithms. In *NeurIPS*, 2017. 1, 8, 12
- [76] Zhewei Yao, Amir Gholami, Qi Lei, Kurt Keutzer, and Michael W. Mahoney. Hessian-based analysis of large batch training and robustness to adversaries. In *Advances in Neural Information Processing Systems*, 2018. 8
- [77] Yaodong Yu, Zitong Yang, Edgar Dobriban, Jacob Steinhardt, and Yi Ma. Understanding generalization in adversarial training via the bias-variance decomposition. *CoRR*, 2021. 1
- [78] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016. 2, 5
- [79] Haichao Zhang and Jianyu Wang. Defense against adversarial attacks using feature scattering-based adversarial training. In *NeurIPS*, 2019. 8
- [80] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *ICML*, 2019. 1, 2, 3, 4, 5, 6, 7, 8
- [81] Jingfeng Zhang, Xilie Xu, Bo Han, Gang Niu, Lizhen Cui, Masashi Sugiyama, and Mohan S. Kankanhalli. Attacks which do not kill training make adversarial learning stronger. In *ICML*, 2020. 1, 6
- [82] Jingfeng Zhang, Jianing Zhu, Gang Niu, Bo Han, Masashi Sugiyama, and Mohan S. Kankanhalli. Geometry-aware instance-reweighted adversarial training. In *ICLR*, 2021. 1
- [83] Linjun Zhang, Zhun Deng, Kenji Kawaguchi, Amirata Ghorbani, and James Zou. How does mixup help with robustness and generalization? In *ICLR*, 2021. 8
- [84] Tianhang Zheng, Changyou Chen, and Kui Ren. Distributionally adversarial attack. In *AAAI*, 2019. 8