DOUBLE DESCENT IN ADVERSARIAL TRAINING: AN IMPLICIT LABEL NOISE PERSPECTIVE

Anonymous authors
Paper under double-blind review

ABSTRACT

Here, we show that the robust overfitting shall be viewed as the early part of an epoch-wise double descent — the robust test error will start to decrease again after training the model for a considerable number of epochs. Inspired by our observations, we further advance the analyses of double descent to understand robust overfitting better. In standard training, double descent has been shown to be a result of label flipping noise. However, this reasoning is not applicable in our setting, since adversarial perturbations are believed not to change the label. Going beyond label flipping noise, we propose to measure the mismatch between the assigned and (unknown) true label distributions, denoted as *implicit label noise*. We show that the traditional labeling of adversarial examples inherited from their clean counterparts will lead to implicit label noise. Towards better labeling, we show that predicted distribution from a classifier, after scaling and interpolation, can provably reduce the implicit label noise under mild assumptions. In light of our analyses, we tailored the training objective accordingly to effectively mitigate the double descent and verified its effectiveness on three benchmark datasets.

1 INTRODUCTION

In adversarial training, a typical phenomenon is that after a certain training epoch, the robust test error will start to increase constantly with further training, despite the robust training error continues to decrease (Rice et al., 2020). This phenomenon, known as *robust overfitting*, is believed to be separated from *double descent* (Belkin et al., 2019) in the literature (Rice et al., 2020).

Here, we find that robust overfitting shall be viewed as the early part of an epoch-wise double descent (Nakkiran et al., 2020). As shown in Figure 1, for relatively large models, the robust test error increases only transiently and will eventually decrease again after a considerable number of training epochs. Moreover, we observe consistent and similar patterns across various training settings (e.g., different sample sizes, optimizers, learning rate schedulers, and neural architectures), which further verifies our intuition on the connection between the robust overfitting and the double descent (more discussions are elaborated in Appendix B.1).

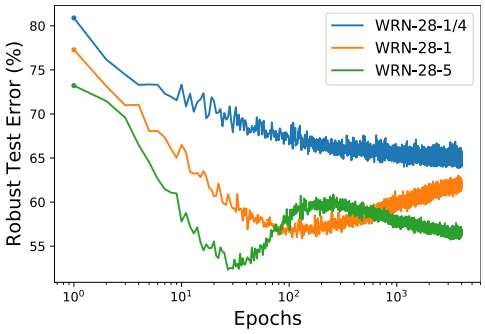

Figure 1: Robust overfitting can be viewed as an early part of the epoch-wise double descent. Here we employ PGD training (Madry et al., 2018) on CIFAR-10 (Krizhevsky, 2009) with Wide ResNet (WRN) (Zagoruyko & Komodakis, 2016) and a fixed learning rate.

Inspired by this observation, we further analyze the epoch-wise double descent to better understand the robust overfitting. In standard training, the double descent has been shown to be related to the label flipping noise (Nakkiran et al., 2020; Yang et al., 2020b). However, adversarial perturbations are believed to not change underlying labels, which is further confirmed empirically from our manual inspection. Therefore, the popular label flipping noise explanation is not applicable here. On the other hand, it is reasonable to believe that adversarial perturbation distorts the label distribution since the examples do become more ambiguous after perturbation. Thus, we focus on the *implicit label noise* which

measures the mismatch between the assigned label distribution and (unknown) true label distribution. Extensive analyses of such implicit label noise reveal the intriguing dependency of robust overfitting on perturbation size (Dong et al., 2021b) and data quality (Dong et al., 2021a).

Guided by our analyses, we design a theoretically-grounded method to mitigate the robust overfitting. The key idea is to resort to an alternative labeling of the adversarial examples. We show that the predictive label distribution of a probabilistic classifier adversarially trained as usual, but after being properly scaled and interpolated, can be utilized as a better labeling of the adversarial examples and provably reduce the implicit label noise. This echoes the recent empirical practice of incorporating knowledge distillation (Hinton et al., 2015) into adversarial training (Chen et al., 2021). While previous works intuitively select fixed scaling and interpolation parameters for knowledge distillation, we show that it is possible to fully unleash the potential of knowledge distillation by automatically determining the set of parameters that maximally reduces the implicit label noise with a strategy similar to confidence calibration (Guo et al., 2017). Such strategy can further mitigate robust overfitting to a minimal amount without additional human tuning effort. Extensive experiments on different datasets, training methods, neural architectures and robustness evaluation metrics verify the effectiveness of our method.

In summary, our findings and contributions are as follows.

- We show that robust overfitting shall be viewed as the early part of an epoch-wise double descent, extending the common belief in adversarial training.
- We show that double descent in adversarial training may originate from the implicit label noise introduced by improper labeling of adversarial examples in adversarial training practice.
- We show an alternative labeling of the adversarial examples can be established to provably reduce the implicit label noise and mitigate the robust overfitting.

The remainder of this paper is organized as follows. In Section 2, we briefly review the existing works that explore robust overfitting and double descent. In Section 3, we theoretically explore the origin of double descent in adversarial training from an implicit label noise perspective. In Section 4, we propose to mitigate the double descent in adversarial training by alternative labeling based on our understanding. Section 5 demonstrates the effectiveness of our method on realistic datasets. Conclusions and further implications are discussed in Section 6.

2 Related Work

Robust overfitting and double descent in adversarial training. Double descent refers to the phenomenon that overfitting by increasing model complexity will eventually improve test set performance (Neyshabur et al., 2017; Belkin et al., 2019). This appears to conflict with the robust overfitting phenomenon in adversarial training, where increasing model complexity by training longer will impair test set performance constantly after a certain point during training. It is thus believed in the literature that double descent and robust overfitting are separate phenomena (Rice et al., 2020).

Towards a more complete understanding of robust overfitting, in this work, we conduct adversarial training for exponentially more epochs than the typical practice. We find that robust overfitting shall be viewed as the early part of an epoch-wise double descent. And increasing the model architecture size, another way to increase the model complexity, can modulate the epoch-wise double descent curve such that either the overfitting curve is shown, or the entire double descent curve is revealed within the same number of training epochs as shown in Figure 1. Therefore, robust overfitting will not go beyond modern generalization theory as an exception and should be adequately explained by the origin of double descent such as label noise.

A recent work also considers a different notion of double descent that is defined with respect to the perturbation size (Yu et al., 2021). Such double descent might be more related to the robustness-accuracy trade-off problem (Papernot et al., 2016; Su et al., 2018; Tsipras et al., 2019; Zhang et al., 2019), rather than the classic picture of double descent based on model complexity.

Understand double descent in adversarial training. In standard training, double descent is often attributed to increased variance, with label noise being a common source. Definitions of label noise vary in literature. Theoretically-grounded analyses of double descent focus on additive label noise (Advani & Saxe, 2020; Mei & Montanari, 2019; Hastie et al., 2019; Belkin et al., 2020; d'Ascoli

et al., 2020), but only applicable to regression problems. Theoretical results on double descent are scarce on classification problems, with a few works introducing noise by randomly masking the feature vector (Deng et al., 2019; Kini & Thrampoulidis, 2020). Analyses of double descent on classification problems are more common in empirical studies, where a typical way to induce double descent is to inject label flipping noise, namely the labels of a random fraction of training examples are flipped to other labels (Nakkiran et al., 2020; Yang et al., 2020b). However, such a definition of label noise cannot properly fit the scenario in adversarial training, where labels are not likely to be flipped due to small adversarial perturbation.

To understand the double descent in adversarial training within the existing picture of double descent, we introduce implicit label noise which is incurred by the mismatch between the assigned label distribution and true label distribution of the adversarial example. Such label noise can be interpreted as an instance-dependent, class dependent label noise in reference to the systematic studies on the taxonomy of label noise (Frénay & Verleysen, 2014; Song et al., 2020), thus can be the origin of double descent in adversarial training.

Mitigate robust overfitting. Robust overfitting hinders the practical deployment of adversarial training methods as the final performance is often sub-optimal. Various regularization methods including classic approaches such as ℓ_1 and ℓ_2 regularization and modern approaches such as cutout (Devries & Taylor, 2017) and mixup (Zhang et al., 2018) have been attempted to tackle robust overfitting, whereas they are shown to perform no better than simply early stopping the training on a validation set (Rice et al., 2020). However, early stopping raises additional concern as the best checkpoint of the robust test accuracy and that of the standard accuracy often do not coincide (Chen et al., 2020), thus inevitably sacrificing the performance on either criterion. Various regularization methods specifically designed for adversarial training are thus being proposed to outperform early stopping, including regularization the flatness of the weight loss landscape (Wu et al., 2020; Stutz et al., 2021), introducing low-curvature activation functions (Singla et al., 2021) and adopting stochastic weight averaging (Izmailov et al., 2018) and knowledge distillation (Hinton et al., 2015) (Chen et al., 2021). These methods are likely to suppress the implicit label noise, with the self-distillation framework (i.e. the teacher shares the same architecture as the student model) introduced by (Chen et al., 2021) as a particular example since introducing teacher's outputs as supervision is almost equivalent to the alternative labeling inspired by our understanding of the origin of label noise in adversarial training.

3 Double Descent from an Implicit Label Noise Perspective

In this section, we present a novel perspective to understand the double descent in adversarial training.

3.1 Preliminaries

Let $\mathcal{X} \subset \mathbb{R}^d$ define the input space equipped with a norm $\|\cdot\| : \mathcal{X} \to \mathbb{R}^+$ and $\mathcal{Y} = \{1, \dots, c\}$ define the label space. Considering a classification problem with a training set $\mathcal{D} = \{(x, y)\}$, where $x \in \mathcal{X}$ is an input and $y \in \mathcal{Y}$ is its assigned label. We denote $p(Y = j|x)$ as the distribution of the assigned labels of x given by a group of annotators. Typically $y = \arg\max_j p(Y = j|x)$. When there is no ambiguity, we will denote $p(Y = j|x)$ as $p(y|x)$ with a slight abuse of notation. Let $y^* \in \mathcal{Y}$ be the true label of x and $p(y^*|x)$ be its distribution, which can be given by a group of experts. Note $y^* = \arg\max_j p(Y^* = j|x)$.

Assumption 3.1. *We assume there is no label noise in $\mathcal{D}$, namely $y = y^*$ and $p(y|x) = p(y^*|x)$.*

Definition 3.1 (Data quality)**.** *Given an example $x \in \mathcal{X}$, we define its data quality as $\max_j p(Y^* = j|x)$, namely the maximum class probability of the true label distribution.*

Assumption 3.2. *We assume the original clean dataset $\mathcal{D}$ contains mostly high-quality data, namely $\max_j p(Y^* = j|x) \approx 1$.*

We now briefly review adversarial training. Let $f(\cdot; \theta) : \mathcal{X} \to \mathcal{Y}$ be a probabilistic classifier and $f(\cdot; \theta)_j$ be its j-th class probability. Adversarial training can be viewed as an data augmentation technique that trains the classifier f on a set of adversarial examples $\mathcal{D}_\delta = \{(x_\delta, y_\delta)\}$, namely

$$\theta^* = \arg\min_\theta \mathbb{E}_{\mathcal{D}_\delta} \, \ell(f(x_\delta; \theta), y_\delta). \tag{1}$$

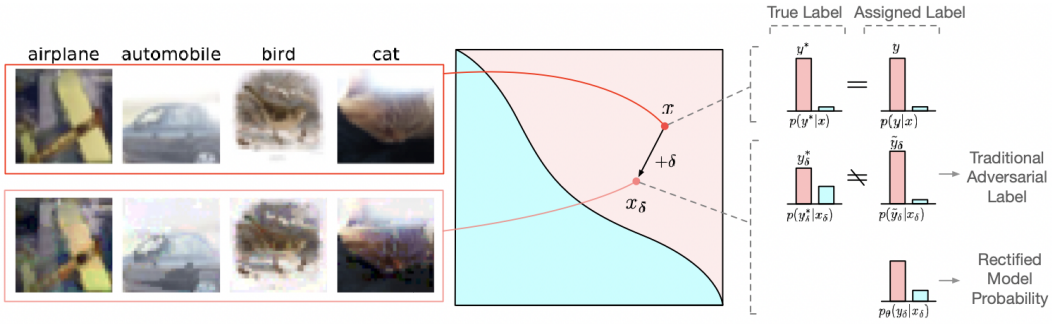

Figure 2: The illustration of the origin of implicit label noise in adversarial training. The traditional adversarial label introduces the implicit label noise by inducing a distribution mismatch between the assigned label distribution and true label distribution of the adversarial example. Our rectified model probability as an alternative labeling can provably reduce this distribution match.

Here, $x_\delta \equiv x + \delta$ is the adversarial example of x produced by the inner maximization, where

$$\delta = \underset{\delta, \|\delta\| \leq \epsilon}{\arg\max}\ \ell(f(x+\delta;\theta), y), \tag{2}$$

and $y_\delta \in \mathcal{Y}$ is the assigned label of x_δ. Finally, we denote $y_\delta^* \in \mathcal{Y}$ as the true label of x_δ.

Remark 3.1. *In adversarial training, it is the common practice that copies the label of a clean example to its adversarial counterpart, namely $y_\delta \leftarrow y$ and $p(y_\delta|x_\delta) \leftarrow p(y|x)$.*

We denote $\tilde{y}_\delta \equiv y$ and its associated distribution as the *traditional adversarial label*. We believe traditional adversarial label is the key problem that induces double descent in adversarial training.

3.2 Traditional adversarial label **does not introduce label flipping noise**

Since label noise is often essential to explain double descent in standard learning for modern neural architectures (Nakkiran et al., 2020; Yang et al., 2020b), we wish to check if the traditional adversarial label produces any label flipping noise, a typical type of label noise, namely if $\tilde{y}_\delta \neq y_\delta^*$ for any adversarial example x_δ. This is equivalent to checking if $y^* \neq y_\delta^*$, since $\tilde{y}_\delta = y$ by the construction of the traditional adversarial label (Remark 3.1) and $y = y^*$ by the construction of our clean dataset (Assumption 3.1). Note that $y^* \neq y_\delta^*$ means the semantics of the adversarial example is distorted significantly such that its true label is now different from the true label of its clean counterpart.

We first visually check the adversarial examples created by the inner maximization in adversarial training. We are specifically interested in those examples with the lowest data quality as they are mostly likely to become label noise after perturbation. We estimate the data quality based on model ensemble as shown in Appendix E.1. One can find that the adversarial counterparts of those low-quality examples still match their original labels, albeit being slightly more ambiguous.

We have also found some evidence from the literature. Yang et al. (2020a) draw the same conclusion by studying the distance between examples in the input space. In the CIFAR-10 training set, they found that the minimum distance between any two examples from two different classes respectively is around $50/255$ in terms of ℓ_∞ norm, while the perturbation radius typically used in adversarial training is around $8/255$, which is significantly smaller. It is thus reasonable to claim that a significant amount of label flipping noise will not be produced by the inner maximization, and subsequently it cannot explain the double descent observed in adversarial training.

3.3 Traditional adversarial label **introduces implicit label noise**

Although traditional adversarial label is unlikely to directly produce label flipping noise, we show that it can produce label noise implicitly.

It has been widely observed that adversarial examples transfer between different classifiers even with distinct architectures (Papernot et al., 2017). Meanwhile, an adversarial example often contains salient characteristics of classes other than the original label from human perspective (Tsipras et al., 2019; Ilyas et al., 2019). Therefore, the true label distribution of an adversarial example is likely to

be different from its clean counterpart, which means $p(y_\delta^*|x_\delta) \neq p(y^*|x)$. Note this will not conflict with $y_\delta^* = y^*$ as discussed in Section 3.2, since the latter only implies $\arg\max_j \; p(Y_\delta^* = j|x_\delta) = \arg\max_j \; p(Y^* = j|x)$.

Now we show the traditional adversarial label is improper in terms of the label distribution. As illustrated in Figure 2, we have $p(\tilde{y}_\delta|x_\delta) = p(y|x)$ by Remark 3.1 and $p(y|x) = p(y^*|x)$ by Assumption 3.1, which together with the effect of adversarial perturbation $p(y_\delta^*|x_\delta) \neq p(y^*|x)$ imply that $p(\tilde{y}_\delta|x_\delta) \neq p(y_\delta^*|x_\delta)$. This means there is a distribution mismatch between the true label distribution and the assigned label distribution of the adversarial example. Such distribution mismatch introduces *implicit label noise* in the training set $\mathcal{D}_\delta$ employed in adversarial training.

Definition 3.2 (Implicit label noise)**.** *Given an example $x \in \mathcal{X}$, we define the implicit label noise as the probability that its assigned label is different from its true label, namely $p(Y \neq Y^*|x)$.*

Remark 3.2. *Implicit label noise is equivalent to instance-dependent and class-dependent label noise, since $p(Y \neq Y^*|x) = 1 - \sum_j p(Y = j|Y^* = j, x)p(Y^* = j|x)$, and $p(Y \neq j|Y^* = j, x) = 1 - p(Y = j|Y^* = j, x)$ is a typical definition of label noise. It can easily seen that if $p(Y \neq Y^*|x) > 0$, $p(Y \neq j|Y^* = j, x) > 0$ for some j.*

Before showing the connection between the distribution mismatch and the *implicit label noise*, we need to quantify the distribution mismatch first. Here we adopt the *total variation distance*.

Definition 3.3 (Total variation (TV) distance)**.** *For two discrete probability distributions $p(y)$ and $p(y')$ where $y, y' \in \mathcal{Y}$, the total variation distance between them can be defined as*

$$\|p(y) - p(y')\|_{TV} = \max_J \left| \sum_{j \in J} p(y = j) - \sum_{j \in J} p(y' = j) \right| \tag{3}$$

We are now ready to present our main result.

Theorem 3.1 (Traditional adversarial label incurs implicit label noise)**.** *The implicit label noise incurred by the distribution mismatch between the traditional adversarial label and the true label of the adversarial example is lower-bounded, namely*

$$p(\tilde{y}_\delta \neq y_\delta^*|x_\delta) \geq \|p(\tilde{y}_\delta|x_\delta) - p(y_\delta^*|x_\delta)\|_{TV} = \|p(y^*|x) - p(y_\delta^*|x_\delta)\|_{TV}. \tag{4}$$

An informal proof can be sketched from a frequentist's view and help the understanding of implicit label noise. Say there are M identical copies of x_δ in the training set $\mathcal{D}_\delta$, with their true labels and traditional adversarial labels distributing according to $p(y_\delta^*|x_\delta)$ and $p(\tilde{y}_\delta|x_\delta)$, respectively. The number of copies that have the same true label and assigned label is $M \sum_j \min\{p(\tilde{Y}_\delta = j|x_\delta), p(Y_\delta^* = j|x_\delta)\}$. The fraction of label noise exists in $\mathcal{D}_\delta$ is thus $1 - \sum_j \min\{p(\tilde{Y}_\delta = j|x_\delta), p(Y_\delta^* = j|x_\delta)\} = \|p(\tilde{y}_\delta|x_\delta) - p(y_\delta^*|x_\delta)\|_{\text{TV}}$ by the definition of the total variation distance.

3.4 Dependence of Implicit Label Noise in Adversarial Training

We now show the implicit label noise in adversarial training depends on the perturbation size and the data quality, which will subsequently affect the double descent curves.

We consider optimal adversarial perturbation for simplicity, which is defined as

$$\delta^* = \arg\max_{\delta, \|\delta\| \leq \varepsilon} \; \ell(f^*(x + \delta; \theta), y), \tag{5}$$

where $f^*(\cdot)$ denotes the optimal probabilistic classifier such that $f^*(x) = p(y^*|x)$. With ℓ being the cross-entropy loss, the above equation can be rewritten as

$$\delta^* = \arg\min_{\delta, \|\delta\| \leq \varepsilon} f^*(x + \delta; \theta)_{y^*}, \tag{6}$$

which implies the optimal perturbation will reduce the probability mass of the true label. With the optimal adversarial perturbation, the minimum implicit label noise can be solved as follows.

Theorem 3.2 (Minimum implicit label noise under optimal perturbation)**.** *Assume the true label distribution $p(y^*|x)$ is locally convex around x and can be asymptotically described as*

$$\|\nabla_x \, p(Y^* = j|x)\| \propto 1 - p(Y^* = j|x), \quad p(Y^* = j|x) \to 1 \tag{7}$$

we have

$$\underline{\min} \, p(\tilde{y}_\delta \neq y^*_\delta | x_\delta) \propto \varepsilon(1 - \max_j p(Y^* = j|x)),$$

where $\underline{\min}$ *means the lower bound of the minimum label noise.*

In Appendix A.1, we show that Assumption (7) holds true for a Gaussian mixture model. The above theorem shows that, when the data quality of the clean example is relatively high, the implicit label noise introduced by adversarial perturbation is proportional to (1) the perturbation radius (2) the data quality. We conduct controlled experiments and empirically verify this correlation in Appendix B.2.

Interestingly, the above analysis implies that if the perturbation does not reduce the probability mass of the true label, it will not cause implicit label noise even with an extremely large perturbation radius. We demonstrate this in detail and empirically verified it for Gaussian noise in Appendix B.3.

3.5 Implicit label noise induces double descent

In standard training, the effect of label noise on double descent has been rigorously studied based upon both analytical settings (Mei & Montanari, 2019; Hastie et al., 2019; Deng et al., 2019; Belkin et al., 2020) and bias-variance analyses (Jacot et al., 2020; Yang et al., 2020b; d'Ascoli et al., 2020). Since implicit label noise is just a special case of label noise (Remark 3.2), and adversarial training can be viewed as standard training on an augmented dataset (Equation (1)), it can be inferred that implicit label noise will cause double descent in adversarial training. To demonstrate this in a straightforward way, in Figure 3 we employ standard training on a dataset augmented by fixed adversarial perturbation and show it can indeed produce double descent.

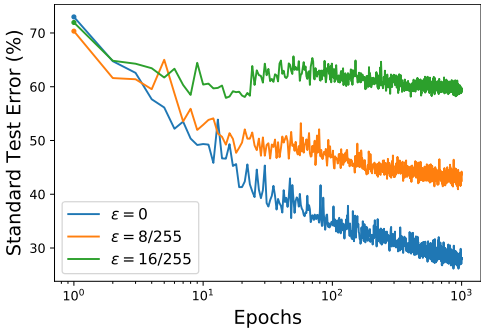

Figure 3: Without explicitly injecting label noise, standard training can also produce double descent if conducted on a dataset augmented by fixed and small adversarial perturbation, which can be properly explained by our implicit label noise perspective. Detailed experiment settings can be found in Appendix F.3.

Since implicit label noise modulates the double descent, and by Theorem 3.2 it depends on the perturbation radius and data quality, the double descent in adversarial training should strongly correlate with the perturbation radius and data quality. Indeed, it has been observed respectively that small perturbation radius will not induce robust overfitting (Dong et al., 2021b), and high-quality data will not induce robust overfitting (Dong et al., 2021a).

4 Mitigate Double Descent in Adversarial Training

In this section, we focus on suppressing the implicit label noise in adversarial training to mitigate double descent from both theoretical and practical perspectives. Since the implicit label noise is incurred by the traditional adversarial label which improperly labels the adversarial example and causes a distribution mismatch between its assigned label and true label, we wish to find an alternative label (distribution) for the adversarial example to reduce such distribution mismatch.

4.1 Rectify model probability to reduce distribution mismatch

We show that it is possible to reduce the distribution mismatch by utilizing the predictive probability of a classifier trained on traditional adversarial labels, which we will refer as the *model probability* for simplicity. We now provide a theoretical guarantee to show that, with *temperature scaling* (Hinton et al., 2015; Guo et al., 2017) enabled in the softmax function, model probability induces a distribution mismatch provably smaller than the traditional adversarial label.

Theorem 4.1 (Model probability induces smaller distribution mismatch than the traditional adversarial label)**.** *Let $f(x;\theta,T)$ denote the predictive probability of a hypothesis classifier scaled by temperature T, namely*

$$f(x;\theta,T)_j = \frac{\exp(z_j/T)}{\sum_j \exp(z_j/T)},$$

where z is the logits of the classifier from x.

Let x_δ be an adversarial example correctly classified by a classifier f, i.e. $\arg\max_j f(x_\delta;\theta)_j = y^*_\delta$, then there exists T, such that

$$\|f(x_\delta;\theta,T) - p(y^*_\delta|x_\delta)\|_{TV} \leq \|p(\tilde{y}_\delta|x_\delta) - p(y^*_\delta|x_\delta)\|_{TV}.$$

One possible way to further reduce the distribution mismatch is to interpolate between the model probability and the traditional adversarial label.

Theorem 4.2 (Interpolation can further reduce the distribution mismatch)**.** *Let x_δ be an adversarial example incorrectly classified by a classifier f, i.e. $\arg\max_j f(x_\delta;\theta,T)_j \neq y^*_\delta$. Assume $\max_j p(Y^*_\delta = j|x_\delta) \geq 1/2$, then there exists an interpolation ratio λ, such that*

$$\|\lambda \cdot f(x_\delta;\theta,T) + (1-\lambda)\cdot p(\tilde{y}_\delta|x_\delta) - p(y^*_\delta|x_\delta)\|_{TV} \leq \|f(x_\delta;\theta,T) - p(y^*_\delta|x_\delta)\|_{TV}.$$

Note the above theorem focus on incorrectly classified examples and thus can be regarded as a complement to Theorem 4.1.

As a summarization, to reduce the distribution mismatch, we propose to use the following distribution as the assigned label of the adversarial example in adversarial training instead of the traditional adversarial label.

$$p_{\theta;T,\lambda}(y_\delta|x_\delta) = \lambda \cdot f(x_\delta;\theta,T) + (1-\lambda)\cdot p(\tilde{y}_\delta|x_\delta), \tag{8}$$

We refer this label distribution as the *Rectified model probability*.

In Appendix D, we show that the optimal hyper-parameters (i.e. T and λ) of almost all training examples concentrate on the same set of values by empirically studying on a synthetic dataset with known true label distribution. Therefore it is possible to find an universal set of hyper-parameters that reduce the distribution mismatch for all adversarial examples.

4.2 Determine the optimal temperature and interpolation ratio

The set of temperature and interpolation ratio in the rectified model probability that maximally reduces the distribution mismatch is not straightforward to find as the true label distribution of the adversarial example is unknown in reality. Fortunately, given a sufficiently large validation dataset as a whole, it is possible to measure the overall distribution mismatch in a frequentist's view without knowing the true label distribution of every single example. A popular metric adopted here is the negative log-likelihood (NLL) loss, which is known as a proper scoring rule (Gneiting & Raftery, 2007) and is also employed in the confidence calibration of deep networks (Guo et al., 2017). By Gibbs's inequality it is easy to show that the NLL loss will only be minimized when the assigned label distribution matches the true label distribution (Hastie et al., 2001), namely

$$-\mathbb{E}_{(x_\delta,y^*_\delta)\in\mathcal{D}^{\text{val}}_\delta} \log p_{\theta;T,\lambda}(Y_\delta = y^*_\delta|x_\delta) \geq -\mathbb{E}_{p(y^*_\delta)\sim\mathcal{D}^{\text{val}}_\delta} p(y^*_\delta|x_\delta)\log p(y^*_\delta|x_\delta). \tag{9}$$

And here y^*_δ is known since the adversarial example should share the same argmax label with its clean counterpart as shown in Section 3.2.

Therefore, we propose to find the optimal T and λ in Equation (8) as

$$T,\lambda = \arg\min_{T,\lambda} -\mathbb{E}_{(x_\delta,y^*_\delta)\in\mathcal{D}^{\text{val}}_\delta} \log p_{\theta;T,\lambda}(Y_\delta = y^*_\delta|x_\delta) \tag{10}$$

One may note that Equation (10) cannot be directly optimized since the traditional adversarial label is only defined on the example in the training set and cannot be simply generalized to the validation set. A reasonable solution is using the nearest neighbour classifier to find the closest traditional adversarial label for every example in the validation set. However, to speed up the optimization we propose to employ the classifier overfitted by the traditional adversarial labels on the training set as an surrogate (see Appendix E.2), which works well in practice. Such process incurs almost no additional computation as we simply obtain the logits of a surrogate classifier on the validation set.

4.3 RECTIFIED MODEL PROBABILITY MITIGATES DOUBLE DESCENT

We now work on a realistic dataset (CIFAR-10) to demonstrate the rectified model probability proposed in Equation (8) can effectively mitigate the robust overfitting, or equivalently the epoch-wise double descent in adversarial training. The outer minimization of adversarial training (Equation (1)) now becomes

$$\theta^* = \arg\min_{\theta} \mathbb{E}_{\mathcal{D}_\delta}\, \ell\left(f(x_\delta;\theta), p_{\theta^{\text{Trad}};T,\lambda}(y_\delta|x_\delta)\right), \tag{11}$$

where θ^{Trad} denotes the parameters of a classifier trained on the traditional adversarial label beforehand. The details of the experimental setting are available in Appendix F.1.

As shown in Figure 4, adversarial training on rectified model probability can best mitigate the robust overfitting when the temperature T and interpolation ratio λ are optimal. Such optimal hyperparameters perfectly aligns with the ones automatically determined by Equation (10).

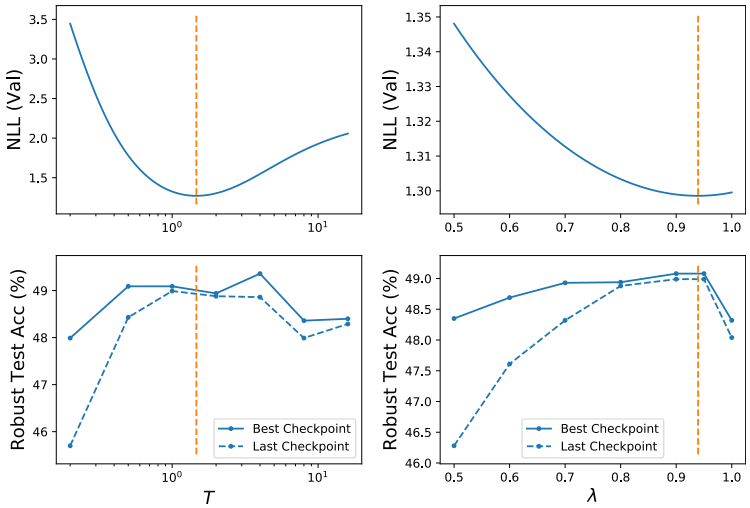

Figure 4: (Upper) NLL loss obtained on the validation set for different T and λ. (Bottom) Robust test accuracy at the best and last checkpoint by adversarial training with the rectified model probability with different T and λ. $\lambda = 0.8$ for grid search on T (Left) and $T = 2$ for grid search on λ (Right). Orange dashed lines in the left and right columns indicate the temperature and the interpolation ratio automatically determined by Equation (10), respectively.

5 EXPERIMENTS

Experiment setup. We consider three datasets, CIFAR-10, CIFAR-100 (Krizhevsky, 2009) and Tiny-ImageNet (Le & Yang, 2015). We consider robustness against ℓ_∞ norm-bounded adversarial attack with perturbation radius $8/255$, and employ AutoAttack (Croce & Hein, 2020) for reliable evaluation. We conduct PGD training on pre-activation ResNet-18 (He et al., 2016) in this section. Appendix C includes results on additional adversarial training methods (e.g., TRADES (Zhang et al., 2019), FGSM (Goodfellow et al., 2015)), neural architectures (e.g., VGG (Simonyan & Zisserman, 2015), WRN) and evaluation metrics (e.g., PGD-1000, Square Attack (Andriushchenko et al., 2020), RayS (Chen & Gu, 2020)). More setup details can be found in Appendix F.

Results & Discussions. We compare our method (*Ours*) with two baselines: regular PGD training on traditional adversarial label (AT), and PGD training combined with self-distillation (KD-AT) with fixed temperature $T = 2$ and interpolation ratio $\lambda = 0.5$ as suggested by Chen et al. (2021).

As shown in Figure 5, our method can effectively mitigate robust overfitting for all datasets, with both standard accuracy (SA) and robust accuracy (RA) constantly increasing throughout training. In Table 1, we measure the difference between the RA at the best checkpoint (Best) and at the last checkpoint (Last) to clearly show the overfitting gap. Our method can reduce the overfitting gap to less than 0.3% for all datasets. One may note that self-distillation with fixed hyper-parameters is

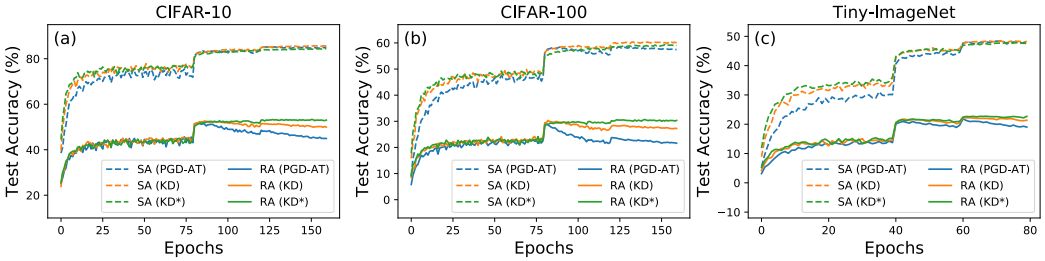

Figure 5: Our method can effectively mitigate robust overfitting for different datasets.

in fact inferior in terms of reducing robust overfitting, while its effectiveness can be significantly improved with the optimal hyper-parameters automatically determined by our method, which further verifies our understanding of robust overfitting.

Compared with self-distillation with fixed hyper-parameters, our method can also boost both RA and SA at the best checkpoint for all datasets.

Dataset	Setting	T	λ	Robust Acc. (%)			Standard Acc. (%)		
				Best	Last	Diff.	Best	Last	Diff.
CIFAR-10	AT	-	-	47.35	41.42	5.93	82.67	84.91	-2.24
	KD-AT	2	0.5	48.76	46.33	2.43	82.89	**85.49**	-2.60
	ours	1.47*	0.8*	**49.05**	**48.80**	**0.25**	**84.26**	84.47	**-0.21**
CIFAR-100	AT	-	-	24.79	19.75	5.04	57.33	57.42	-0.09
	KD-AT	2	0.5	25.77	23.58	2.19	57.24	**60.04**	-2.80
	ours	1.53*	0.83*	**26.36**	**26.24**	**0.12**	**58.80**	59.05	**-0.25**
Tiny-ImageNet	AT	-	-	17.20	15.40	1.80	47.72	47.62	0.10
	KD-AT	2	0.5	17.86	17.18	0.68	**47.73**	**48.28**	-0.55
	ours	1.23*	0.85*	**18.29**	**18.39**	**-0.10**	47.46	47.56	**-0.10**

Table 1: Performance of our method on different datasets. * denotes the hyper-parameter automatically determined by our method.

6 Conclusion and Discussions

In this paper, we extend the understanding of robust overfitting and show that it is the early part of an epoch-wise double descent in adversarial training. Our further analyses show that the double descent may originate from the implicit label noise introduced by the improper labeling of the adversarial examples. Based on our understanding, we propose an alternative labeling of adversarial examples by rectifying model probability. Our method can effectively mitigate robust overfitting without any manual hyper-parameter tuning. In the future, we also plan to explore if it is possible to rectify the model probability on the fly during the adversarial training for a better efficiency.

By showing robust overfitting and epoch-wise double descent are in fact the same phenomena, our work eases the difficulty in understanding overfitting in modern generalization theory. Existing theories on double descent can be readily applied to adversarially robust learning setting, and potentially consolidate the theoretical grounding of robust learning.

Beyond the double descent in adversarial training, our definition of implicit label noise from a distribution mismatch perspective may also exist in a variety of scenarios, especially in real-world tasks where labeling is often complicated and ambiguous. Our method to rectify such label noise should thus be applicable as well. This would be an interesting future work.

Finally, on the practical side, our understanding of robust overfitting opens new opportunities to advance adversarial training practice. Better labeling of the adversarial examples may be still in demand to further reduce robust overfitting.

REPRODUCIBLITY STATEMENT

We conduct experiments on public benchmark. We will release implementations for all methods and scripts for all experiments on GitHub, under the Apache-2.0 license.

ETHIC STATEMENT

In this paper, we explore the double descent phenomenon in adversarial robust learning from the implicit label noise perspective. We experiment on three benchmark datasets that are publicly available. Our analyses offer more in-depth understandings about adversarial training and can better improve the robustness of neural network models. Therefore, we believe our work is ethically on the right side of spectrum and has no potential for misuse, and cannot harm any vulnerable population.

REFERENCES

Madhu S. Advani and Andrew M. Saxe. High-dimensional dynamics of generalization error in neural networks. *Neural Networks*, 132:428 – 446, 2020.

Maksym Andriushchenko, F. Croce, Nicolas Flammarion, and M. Hein. Square attack: a query-efficient black-box adversarial attack via random search. *ArXiv*, abs/1912.00049, 2020.

Mikhail Belkin, Daniel J. Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116:15849 – 15854, 2019.

Mikhail Belkin, Daniel J. Hsu, and J. Xu. Two models of double descent for weak features. *SIAM J. Math. Data Sci.*, 2:1167–1180, 2020.

Y. Carmon, Aditi Raghunathan, L. Schmidt, Percy Liang, and John C. Duchi. Unlabeled data improves adversarial robustness. *ArXiv*, abs/1905.13736, 2019.

J. Chen and Quanquan Gu. Rays: A ray searching method for hard-label adversarial attack. *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020.

Tianlong Chen, Sijia Liu, Shiyu Chang, Yu Cheng, Lisa Amini, and Zhangyang Wang. Adversarial robustness: From self-supervised pre-training to fine-tuning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 696–705, 2020.

Tianlong Chen, Zhenyu Zhang, Sijia Liu, Shiyu Chang, and Zhangyang Wang. Robust overfitting may be mitigated by properly learned smoothening. In *International Conference on Learning Representations*, volume 1, 2021.

F. Croce and M. Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *ICML*, 2020.

Stéphane d'Ascoli, Maria Refinetti, G. Biroli, and F. Krzakala. Double trouble in double descent : Bias and variance(s) in the lazy regime. *ArXiv*, abs/2003.01054, 2020.

Jia Deng, W. Dong, R. Socher, L. Li, K. Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.

Zeyu Deng, A. Kammoun, and Christos Thrampoulidis. A model of double descent for high-dimensional binary linear classification. *ArXiv*, abs/1911.05822, 2019.

Terrance Devries and Graham W. Taylor. Improved regularization of convolutional neural networks with cutout. *ArXiv*, abs/1708.04552, 2017.

Chengyu Dong, Liyuan Liu, and Jingbo Shang. Data profiling for adversarial training: On the ruin of problematic data. *ArXiv*, abs/2102.07437, 2021a.

Yinpeng Dong, Ke Xu, Xiao Yang, Tianyu Pang, Zhijie Deng, Hang Su, and Jun Zhu. Exploring memorization in adversarial training. *ArXiv*, abs/2106.01606, 2021b.

Benoît Frénay and M. Verleysen. Classification in the presence of label noise: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 25:845–869, 2014.

T. Gneiting and A. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102:359 – 378, 2007.

Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *CoRR*, abs/1412.6572, 2015.

Sven Gowal, Chongli Qin, Jonathan Uesato, Timothy A. Mann, and P. Kohli. Uncovering the limits of adversarial training against norm-bounded adversarial examples. *ArXiv*, abs/2010.03593, 2020.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. *ArXiv*, abs/1706.04599, 2017.

T. Hastie, A. Montanari, S. Rosset, and R. Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *ArXiv*, abs/1903.08560, 2019.

Trevor J. Hastie, Robert Tibshirani, and Jerome H. Friedman. The elements of statistical learning. 2001.

Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. *ArXiv*, abs/1603.05027, 2016.

Dan Hendrycks and Thomas G. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *ArXiv*, abs/1903.12261, 2019.

Geoffrey E. Hinton, Oriol Vinyals, and J. Dean. Distilling the knowledge in a neural network. *ArXiv*, abs/1503.02531, 2015.

Andrew Ilyas, Shibani Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry. Adversarial examples are not bugs, they are features. In *NeurIPS*, 2019.

Pavel Izmailov, Dmitrii Podoprikhin, T. Garipov, D. Vetrov, and A. Wilson. Averaging weights leads to wider optima and better generalization. *ArXiv*, abs/1803.05407, 2018.

Arthur Jacot, Berfin Simsek, F. Spadaro, Clément Hongler, and F. Gabriel. Implicit regularization of random feature models. In *ICML*, 2020.

Ganesh R. Kini and Christos Thrampoulidis. Analytic study of double descent in binary classification: The impact of loss. *2020 IEEE International Symposium on Information Theory (ISIT)*, pp. 2527–2532, 2020.

A. Krizhevsky. Learning multiple layers of features from tiny images. 2009.

Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. 2015.

I. Loshchilov and F. Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv: Learning*, 2017.

A. Madry, Aleksandar Makelov, L. Schmidt, D. Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *ArXiv*, abs/1706.06083, 2018.

Song Mei and A. Montanari. The generalization error of random features regression: Precise asymptotics and double descent curve. *arXiv: Statistics Theory*, 2019.

Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, B. Barak, and Ilya Sutskever. Deep double descent: Where bigger models and more data hurt. *ArXiv*, abs/1912.02292, 2020.

Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nathan Srebro. Exploring generalization in deep learning. In *NIPS*, 2017.

Nicolas Papernot, P. McDaniel, Arunesh Sinha, and Michael P. Wellman. Towards the science of security and privacy in machine learning. *ArXiv*, abs/1611.03814, 2016.

Nicolas Papernot, P. Mcdaniel, I. Goodfellow, S. Jha, Z. Y. Celik, and A. Swami. Practical black-box attacks against machine learning. *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, 2017.

Leslie Rice, Eric Wong, and J. Z. Kolter. Overfitting in adversarially robust deep learning. *ArXiv*, abs/2002.11569, 2020.

K. Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2015.

Vasu Singla, Sahil Singla, David Jacobs, and Soheil Feizi. Low curvature activations reduce overfitting in adversarial training. *ArXiv*, abs/2102.07861, 2021.

Leslie N. Smith. Cyclical learning rates for training neural networks. *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 464–472, 2017.

Hwanjun Song, Minseok Kim, Dongmin Park, and J. Lee. Learning from noisy labels with deep neural networks: A survey. *ArXiv*, abs/2007.08199, 2020.

David Stutz, Matthias Hein, and B. Schiele. Relating adversarially robust generalization to flat minima. *ArXiv*, abs/2104.04448, 2021.

D. Su, Huan Zhang, H. Chen, Jinfeng Yi, P. 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.

D. Tsipras, Shibani Santurkar, L. Engstrom, A. Turner, and A. Madry. Robustness may be at odds with accuracy. *arXiv: Machine Learning*, 2019.

Jonathan Uesato, Jean-Baptiste Alayrac, Po-Sen Huang, Robert Stanforth, Alhussein Fawzi, and P. Kohli. Are labels required for improving adversarial robustness? *ArXiv*, abs/1905.13725, 2019.

Dongxian Wu, Shutao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. *arXiv: Learning*, 2020.

Y. Yang, Cyrus Rashtchian, Hongyang Zhang, R. Salakhutdinov, and K. Chaudhuri. A closer look at accuracy vs. robustness. *arXiv: Learning*, 2020a.

Zitong Yang, Yaodong Yu, Chong You, J. Steinhardt, and Y. Ma. Rethinking bias-variance trade-off for generalization of neural networks. In *ICML*, 2020b.

Yaodong Yu, Zitong Yang, E. Dobriban, J. Steinhardt, and Yi Ma. Understanding generalization in adversarial training via the bias-variance decomposition. *ArXiv*, abs/2103.09947, 2021.

Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *ArXiv*, abs/1605.07146, 2016.

Hongyang Zhang, Yaodong Yu, J. Jiao, E. Xing, L. Ghaoui, and Michael I. Jordan. Theoretically principled trade-off between robustness and accuracy. *ArXiv*, abs/1901.08573, 2019.

Hongyi Zhang, Moustapha Cissé, Yann Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *ArXiv*, abs/1710.09412, 2018.