

FAIR-TAT: Improving Model Fairness Using Targeted Adversarial Training

Tejaswini Medi¹, Steffen Jung^{1,2}, Margret Keuper^{1,2}

¹University of Mannheim, Germany ²MPI for Informatics, Saarland Informatics Campus, Germany
tejaswini.medi@uni-mannheim.de

Abstract

Deep neural networks are susceptible to adversarial attacks and common corruptions, which undermine their robustness. In order to enhance model resilience against such challenges, Adversarial Training (AT) has emerged as a prominent solution. Nevertheless, adversarial robustness is often attained at the expense of model fairness during AT, i.e., disparity in class-wise robustness of the model. While distinctive classes become more robust towards such adversaries, hard to detect classes suffer. Recently, research has focused on improving model fairness specifically for perturbed images, overlooking the accuracy of the most likely non-perturbed data. Additionally, despite their robustness against the adversaries encountered during model training, state-of-the-art adversarial trained models have difficulty maintaining robustness and fairness when confronted with diverse adversarial threats or common corruptions. In this work, we address the above concerns by introducing a novel approach called Fair Targeted Adversarial Training (FAIR-TAT). We show that using targeted adversarial attacks for adversarial training (instead of untargeted attacks) can allow for more favorable trade-offs with respect to adversarial fairness. Empirical results validate the efficacy of our approach.

1. Introduction

Deep neural networks have led to tremendous success in a wide range of computer vision applications [12, 18, 22]. One of the most relevant remaining challenges is their lack of robustness towards small domain shifts [32], common corruptions [19, 21] and adversarially perturbed examples [3, 11, 13, 24]. This limits their usage in safety-critical applications [28]. Adversarial examples are artificially crafted samples, which do not alter the semantic meaning of natural samples. However, they add a small amount of noise that alters the model’s prediction. As a prominent defense mechanism to such adversaries, models are trained using adversarial training (AT) algorithms [30, 42, 49, 52], where adversarial samples along with clean samples are utilized

during model training to build a more robust classifier. The most commonly used adversarial examples during AT are generated through Fast Gradient Sign Method (FGSM) [13] and Projected Gradient Descent (PGD) [30]. These methods are overall efficient to compute and have shown to improve the overall model robustness [30, 42, 49].

Despite improving the overall model robustness, adversarial training has side effects: it usually reduces model accuracy on clean samples, and, although accuracy on perturbed samples improves, there remains a noticeable accuracy gap between clean and perturbed samples [40, 41, 46, 49]. A further significant observation is the impact of AT on model fairness [6, 44, 45], where a lack of fairness is measured by the disparity in class-wise performance on clean data and under attack. Specifically, adversarially trained models tend to exhibit stronger robustness on distinctive classes but weak performance on harder to predict classes. Thus, classes that are difficult to classify in clean data are even more susceptible to exploitation by attackers. To avoid class-wise safety issues and a false sense of security for models with high average robustness, there is a strong need to mitigate the model discrimination against particular groups of classes and focus on improving performance on harder classes. To provide a simple, not politically motivated example, consider a model that tends to confuse dark feline, say black panthers, with brown bears but is good at distinguishing polar bears and white tigers. All classes will benefit from AT to some extent, when attacked. Yet, the benefit is small for black panthers and brown bears while it is large for tigers and polar bears. At the same time, the clean accuracy will drop for all classes, even for those that benefit from AT on perturbed data.

Some works try to address class-wise biases and robustness differences by imposing fairness constraints during AT. This can involve adjusting the weights of easy and hard classes [45] or changing the perturbation margin of attacks based on the class-wise performance of the model during AT [44, 45]. Recent works on model fairness [26, 29, 38, 43, 45, 48, 52] focus on improving the robustness of the worst-performing classes while preserving overall model robustness. These works typically train using *untargeted* adversaries, i.e. adversaries that maximize

the model’s loss irrespective of the new predicted label (target class). In this work, we shift our focus on vulnerable classes (hard classes) and these classes are targeted more often than the stronger ones (easy classes) by generating more *targeted* adversaries towards hard classes. This approach helps to defend against perturbations targeting hard classes, and therefore reduces existing class-wise biases.

Our proposed AT framework is called Fair Targeted Adversarial Training (FAIR-TAT). Unlike conventional AT, which focuses on training with *untargeted* adversaries, FAIR-TAT focuses on *targeted* adversaries. This helps to better understand and steer the class-to-class

biases by adjusting the training objectives of each class during AT. We first examine class-to-class biases using standard metrics like "robust accuracy" and "class-wise false positive scores". Based on the analysis using these metrics, we introduce a method that dynamically adjusts the adversarial training configuration during training. Specifically, we sample the target classes for generating targeted adversaries during training by leveraging the dynamically monitored biases. As a result, models trained using FAIR-TAT show improved accuracy on hard classes for clean data as well as on perturbed data, i.e. their model fairness is improved when compared to classical AT. This finding also extends from adversarial attacks to common corruptions [19].

The contributions of our work can be summarized as follows:

- We empirically examine class-specific vulnerabilities to misclassifications by analyzing class-wise false positives in adversarially trained models. Our findings show that these false positives highlight the most confused classes during training. Based on this, we prioritize these confused classes through generating targeted adversaries of these classes more often during AT and thus learn the class-specific biases.
- We propose a framework called Fair Targeted Adversarial Training (FAIR-TAT), which uses targeted adversaries and adapts the training setup based on class-specific characteristics. Target class selection is guided by the distribution of class-wise false positives seen during instance-wise training. This approach ensures that samples are perturbed to resemble confused classes more frequently, improving class-wise accuracy balance at inference.
- FAIR-TAT can be combined with existing methods to enhance class-wise fairness. We show that FAIR-TAT surpasses state-of-the-art methods in fairness, not only against its trained adversaries but also against other adversaries and common corruptions [19], while maintaining overall robustness.

2. Related Work

Adversarial Robustness. Adversarial robustness is a crucial concept for assessing the resilience of deep neural networks against adversarial attacks [3, 11, 13, 24]. Significant progress has been made in evaluating robustness across diverse adversarially trained model architectures thoroughly [8, 15], while some studies have aimed to unify research efforts on architectural design choices and their influence on robustness [20]. Furthermore, advancements in adversarial robustness have been achieved across tasks such as image classification, restoration, and segmentation by enhancing or redesigning robust architectures [1, 2, 16, 17, 20, 27]. Despite these efforts, which focus on improving overall robust accuracy, consistent variations in robustness across specific classes have been observed. These findings highlight the nuanced challenges in achieving uniformly robust performance across all categories.

Adversarial Training. Models that are not trained with adversarial defenses are typically only robust to low-budget attacks, if at all. Adversarial training (AT) [5, 13, 30, 47] is a remedy to this problem, as it trains the model on adversarial samples optimized during training, effectively making out-of-domain attacks become in-domain samples. However, this approach can result in over-fitting to attacks used during training. Early stopping [37] and the addition of external (synthetic) data [14, 36, 43] have been proposed as effective solutions to address this problem. Prior works have analyzed AT from different perspectives including robust optimization [42], robust generalization [35], and diverse training strategies [34, 42, 50].

Adversarial Fairness Problem. Adversarial Fairness in machine learning is a growing concern, as models often favor specific groups of classes while underperforming on others after being trained adversarially. This discrimination undermines the reliability of machine learning models. Recent studies have investigated class biases in terms of model accuracy and robustness, advocating for fair training processes that classify all classes equally well [31, 39, 44, 45]. These works also show that while AT can improve overall robustness, it often ignores discrepancies in the robustness of individual classes, amplifying unbalance in class-wise robustness. Data distribution and adversarial learning algorithms can contribute to these class-wise discrimination problems, making hard classes more vulnerable to attacks.

To address unfairness, several methods have emerged. Fair Robust Learning (FRL) [45] adjusts perturbation margins and class weights based on whether fairness constraints are violated, though this approach decreases overall accuracy. Balanced Adversarial Training (BAT) [38] aims to balance fairness between source and target classes. Class-wise

Calibrated Fair Adversarial Training (CFA) [44] addresses fairness by dynamically customizing adversarial configurations for different classes and modifying weight averaging techniques. All these methods reweigh classes based on their vulnerability in terms of robustness during the optimization process. Worst Class Adversarial Training (WAT) [26] uses a no-regret dynamic algorithm to improve worst-class accuracy. More recent work, DAFA [25], considers class similarity and assigns distinct adversarial margins and loss weights to each class, adjusting them to encourage a trade-off in robustness among similar classes.

However, there exist classes which exhibit higher robustness performance but also frequently tend to be confused with other classes. Therefore, in this work, we focus on customizing the training configuration of Adversarial Training (AT) based on class-wise confusion rates, which are evaluated using class-wise false positive scores.

3. Preliminaries

AT is defined as min-max optimization problem. Given a model f with parameters θ , loss function L , and training data distribution D , the training algorithm tries to minimize the loss where the adversary aims to maximize the loss within a neighborhood, where the perturbation bound is $\mathcal{B}(x, \epsilon) = \{x' : \|x' - x\|_p \leq \epsilon\}$ and ϵ is the perturbation margin. Equation 1 formalizes the AT objective as a min-max optimization problem, where x' represents the perturbed samples and x represents the clean or natural samples.

$$\min_{\theta} \mathbb{E}_{(x,y) \in D} \left[\max_{\|x' - x\|_p \leq \epsilon} L(f_{\theta}(x'), y) \right] \quad (1)$$

White-box attacks, i.e. attacks with full access to the model and its weights, allow to efficiently find adversaries and are therefore commonly used for AT. FGSM [13] creates adversarial examples in a single step by updating the randomly initialized image perturbation in the direction of the maximal loss increase. Specifically, it computes for an input image x and its label y an adversary x' according to Equation 2, where small values of ϵ ensure that perturbations are imperceptibly small:

$$x' = x + \epsilon \cdot \text{sign}(\nabla_x L(f_{\theta}(x), y)). \quad (2)$$

PGD [30] finds adversaries within the perturbation bound solving the inner maximization problem defined in Equation 1. It employs a multi-step scheme to maximize the inner part of the saddle point formulation. Equation 3 defines the PGD (untargeted) attack, where Π represents the projection function, and α controls the gradient ascent step size:

$$x^{t+1} = \Pi_{\mathcal{B}(x,\epsilon)}(x^t + \alpha \cdot \text{sign}(\nabla_x L(f_{\theta}(x^t), y))). \quad (3)$$

Different from the above adversarial training works, where untargeted white-box attacks are used to find the adversaries, we focus on *targeted adversaries* to solve the

adversarial model fairness problem. In our AT approach, we utilize PGD (targeted) attack [30] to find adversaries within the perturbation bound. Equation 4 defines the PGD (targeted) attack, where y_t defines the target labels:

$$x^{t+1} = \Pi_{\mathcal{B}(x,\epsilon)}(x^t + \alpha \cdot \text{sign}(\nabla_x L(f_{\theta}(x^t), y_t))). \quad (4)$$

Other additional methods, such as those proposed by [7, 33], are often used as probes to determine a model's accuracy under adversarial attack, referred to as the model's *robust accuracy*. Adversarial model fairness in classification is typically evaluated based on the class-wise robust accuracy performance of the classifier. We also evaluate our approach using other standard adversarial attacks like AutoAttack, Squares [4, 10].

Notation. Following the notation in [44], we consider a function $f_{\theta} : X \rightarrow Y$ parameterized by θ representing the K -class classification problem with $Y = \{1, 2, \dots, K\}$. We consider an example x from X and respective label y from Y , and perturb x within the ℓ_p -norm perturbation bound $B(x, \epsilon) = \{x' \mid \|x' - x\|_p \leq \epsilon\}$. Then the clean and robust accuracy of the model f_{θ} are defined as $A(f_{\theta})$ and $R(f_{\theta})$:

$$A(f_{\theta}) = \mathbb{E}_{(x,y) \sim D} [\chi(f_{\theta}(x) = y)], \quad (5)$$

$$R(f_{\theta}) = \mathbb{E}_{(x,y) \sim D} [\chi(\forall x' \in B(x, \epsilon), f_{\theta}(x') = y)], \quad (6)$$

where $\chi(\cdot)$ is the indicator function, equaling 1 if the condition inside is true and 0 otherwise.

Measuring Class-Wise Biases. Different metrics can be employed to assess the class biases of classification models. A reliable aggregate measure is the **Class-Wise Accuracy** (C_{acc}), computed for each class c_j with $j \in \{1, \dots, K\}$, which has been predominantly used in previous works [6, 39], and can be defined from samples of a validation set $\{x_i\}_{i=1}^N$ of size N with labels $y_i \in \{c_j\}_{j=1}^K$ as:

$$C_{acc}(c_j) = \frac{|\{x_i \mid f_{\theta}(x_i) = c_j, y_i = c_j\}|}{N} + \frac{|\{x_i \mid f_{\theta}(x_i) \neq c_j, y_i \neq c_j\}|}{N}, \quad (7)$$

where we consider all samples x_i with label y_i that are correctly classified as class c_j by the model f_{θ} and all samples x_i with label y_i that are correctly classified as not being of class c_j . Previous approaches such as [6, 39, 44] predominantly use this metric to assess class-wise biases and further to calibrate adversarial training by scaling ϵ based on their class-wise robustness performance for each individual class during AT to achieve robust fairness [44].

As a further informative measure on class-wise biases, we consider the **Class False Positive Score** (C_{FPS}) to assess the confusion rates of each class c_j , i.e., how frequently

samples from other classes are incorrectly classified as c_j . We use this measure to evaluate how the other group of classes are attracted to a specific class during AT. While a high C_{FPS} can also manifest in a decrease in class-wise accuracy, the imbalance it induces can hardly be addressed in an untargeted training setting. It therefore measures the potential of leveraging targeted adversarial training for model fairness. To calculate the C_{FPS} for a specific class, we calculate the number of misclassifications where samples from other classes, i.e. samples x_i from the validation set $\{x_i\}_{i=1}^N$ with labels $y_i \in \{c_j\}_{j=1}^K$ are incorrectly classified by model f_θ as this particular class c_j , i.e. the cardinality of $\{x_i | f_\theta(x_i) = c_j, y_i \neq c_j\}$. We then normalize this count by the total number of misclassifications across all classes:

$$C_{FPS}(c_j) = \frac{|\{x_i | f_\theta(x_i) = c_j, y_i \neq c_j\}|}{|\{x_i | f_\theta(x_i) \neq y_i\}|}. \quad (8)$$

A higher C_{FPS} for a class c_j indicates that c_j is frequently assigned by the network even for samples of different classes. Regarding the latent space, one can hypothesize that a class with a high C_{FPS} occupies a relatively large, under-determined region.

4. Empirical Analysis on C_{FPS}




Figure 1. Class False Positive Scores (C_{FPS}) on the CIFAR-10 dataset using adversarially trained models.

To better understand class vulnerability towards misclassifications, we analyze the C_{FPS} to identify classes frequently misassigned by the model. Utilizing this information, we aim to defend against perturbations leading to incorrect predictions of these classes by incorporating adversarial examples targeted towards these classes during AT.

We conduct an experiment by evaluating C_{FPS} on the CIFAR-10 dataset, employing robust models trained adversarially that are sourced from a standardized adversarial

robustness benchmark [8, 15]. Figure 1 illustrates the C_{FPS} scores for CIFAR-10 classes. A clear trend emerges in the C_{FPS} scores across different adversarially trained architectures. The horizontal axis in Figure 1 represent the 10 classes of CIFAR-10 in order. Notably, class "cat", is frequently wrongly assigned to samples from other classes. Similarly, class "airplane", although considered one of the easier classes due to its relatively high robust accuracy, also shows significant misclassification rates. This indicates a need for a deeper understanding of the decision boundaries of these classes when perturbed.

To address the high confusion rates of specific classes, our work uses C_{FPS} to model a prior multinomial distribution of classes. This allows us to sample highly confused classes more frequently during *targeted* adversarial training, where we employ targeted adversaries towards these classes to harden the model against class specific perturbations.

5. FAIR-TAT

Based on the above empirical findings, prioritizing often confused classes during adversarial training can enhance the model’s fairness by directing more adversarial samples towards these specific classes. This allows to reduce target-specific confusion. Therefore, we incorporate targeted adversaries towards these classes more often during AT. When generating targeted adversaries, the choice of which targets to pick for training is crucial. **Should we target all classes equally to promote a uniform learning across all classes during AT? Or should we focus on specific target classes that are hard classes to reduce already existing class-wise imbalances?** This decision shapes the fairness and effectiveness of the training process. The optimal choice depends on the desired outcomes and data characteristics. In this work, we employ a score-based target sampling through evaluation of the false positive scores for individual classes for every instance during AT. Our aim is thus to harden the worst performing classes, by allowing to better define their decision boundaries, not only when being used as source classes but also when defining the target class of an adversarial attack. Please refer to Figure 2 for a visualization of our main idea of FAIR-TAT.

Sampling of Targets. As discussed in Section 3, a high value of C_{FPS} for a specific class indicates frequent misclassifications as other classes. This vulnerability presents an opportunity to enhance model fairness by refining the decision boundary for this specific class. By increasing the frequency of sampling for this class during training, we aim to improve the model’s ability to correctly classify it, thereby reducing misclassifications towards it.

To achieve this, we evaluate C_{FPS} for each instance during training and utilize these class-wise false positive




Figure 2. We visualize the main idea of FAIR-TAT for a hypothetical three classes scenario in a two-dimensional latent space. Here, both classes red (circles) and blue (squares) are hard to distinguish from each other (and hence data points in the latent space are close to each other and the decision boundary), while green (triangles) is a distinctive class where data points are far from the decision boundary. When perturbed via an adversarial attack, data points move (and cross) the decision boundary along indicated paths (arrows). **(left)** During untargeted adversarial training, data points cross the decision boundary that is closest to them, which means no data point is perturbed towards the green class. Hence, the resulting new decision boundary (dashed red) is pushed away from the green class, which further improves its robustness and weakens the robustness of classes red and blue. **(right)** During targeted adversarial training, targets are sampled based on C_{FPS} and therefore some red and blue data points are also targeted towards class green. This leads to an overall more fair decision boundary (dashed blue).

scores as weights to model a prior multi-nominal distribution $\psi(0, K - 1)$. This C_{FPS} based prior distribution guides the sampling of targets $y_t \sim \psi(0, K - 1)$ for the next training instance, where K is number of unique classes. Consequently, classes with high C_{FPS} are sampled more often, indicating more weight on hard challenging classes throughout the training. We ensure that sampled targets do not coincide with ground truth classes by setting their respective sampling probability to zero. Research suggests that datasets like CIFAR-10 and CIFAR-100 contain classes that are challenging to distinguish due to closer semantic similarities among the classes. By prioritizing these challenging or hard classes through sampling and by learning to defend against perturbations leading to misclassification towards these classes, we aim to improve the model’s understanding of the semantic differences between the samples and improve fairness. We have also considered uniform sampling of targets $y_t \sim \mu(0, K - 1)$ to showcase the effectiveness of target sampling using prior based on C_{FPS} . The results with respect to uniform sampling of targets are provided in the supplementary (see Section A.2).

FAIR-TAT Framework. The overall FAIR-TAT framework is shown in [Algorithm 1](#). Suppose for each adversarial training instance, assuming there are $y \in Y$ ground truth classes and $y_t \in Y$ target classes, the training robust accuracy of the specific k -th class during the last training interval

Algorithm 1 FAIR-TAT

```

1: Input: A classifier  $f_\theta(\cdot)$ ; Training dataset  $D$  with  $N$  samples; Ground truth classes  $Y$ ; Unique number of classes  $K$ , Prior distribution based on  $C_{FPS}$   $\psi(0, K - 1)$ ; Uniform distribution  $\mu(0, K - 1)$ , where  $K$  is total number of classes; Batch size  $B$ ; Weight average decay rate  $\alpha$ ; Learning rate  $\eta$ ; Training epochs  $S$ ; Perturbation margin  $\epsilon$ ; Adversarial loss  $\ell$ ; Classification loss  $L$ ; Hyperparameter to scale base perturbation  $\lambda_1$ 
2: Output: FAIR robust classifier  $\hat{f}_{\bar{\theta}}(\cdot)$ 
3: /* Initialize parameters */
4:  $\theta \leftarrow \theta_0, \bar{\theta} \leftarrow \theta$ ;
5: Split  $D = D_{\text{train}} \cup D_{\text{valid}}$ ;
6: for  $y \in Y$  do /* Initialize  $\epsilon_y$  */
7:  $\epsilon_y \leftarrow \epsilon$ ;
8: end for
9: for every epoch  $S \leftarrow 1, 2, 3, \dots, N$  do
10: for every minibatch  $(x, y)$  in  $D_{\text{train}}$  do
11: if same based on  $C_{FPS}$  prior then
12: Sample targets:  $y_t \sim \psi(0, K - 1)$ ;
13: else
14: Sample targets uniformly:  $y_t \sim \mu(0, K - 1)$ ;
15: end if
16: if Any of  $y_t == y$  then
17: Resample targets again with respective target sampling
18: end if
19: Compute optimized targeted adversarial perturbation  $\delta^*$ 
20:  $\delta^* = \arg \min_{\|\delta\| \leq \epsilon} \ell(f_\theta(x + \delta), y_{\text{target}})$ ;
21:  $x' = x + \delta^*$ ;
22:  $\theta \leftarrow \theta - \eta \Delta_\theta L_\theta(x', y)$ ;
23: for  $y \in Y$  and  $y_t \in Y$  do
24: /* Update  $R_k, C_{FPS-k}$  and  $\epsilon_k$  */
25:  $R_k \leftarrow \text{Train-Accuracy}(f_\theta, S)$ ;
26:  $C_{FPS-k} \leftarrow C_{FPS}(f_\theta, S)$ ;
27:  $\epsilon_k \leftarrow (\lambda_1 + r_k)\epsilon$ ;
28: end for
29: end for
30: end for
31: return  $\hat{f}_{\bar{\theta}}$ ;

```

is denoted by R_k . Similarly, C_{FPS-k} represents the class false positive score of the k -th class during the last training interval. We aim to leverage both these properties in our approach. We utilize the class false positive scores to model a multinomial distribution which acts as prior for sampling the targets during our targeted adversarial scenario. We use robust accuracy for customizing the perturbation margin for each class as showcased in [44], that various classes exhibit distinct preferences for perturbation margins during adversarial training, a critical factor in improving fairness. Having smaller ϵ for hard classes than easy classes improves robustness and fairness of the model. We use the update formula from [44] for scaling the ϵ , which is mentioned in line 27 of [Algorithm 1](#). During inference, we use the same ϵ for every class to compare the results with other baselines that focus on fairness. During the training phase, the customized perturbation margin ϵ_k converges to the optimal perturbation margin for each class.

6. Experiments

In this section, we demonstrate the effectiveness of our approach, FAIR-TAT, in enhancing the clean accuracy while maintaining the overall adversarial robustness, along with improving class-wise fairness towards clean and robust sam-

Table 1. Comparison of PGD evaluations of FAIR-TAT framework with other methods focused on model fairness using PRN-18 on CIFAR-10 and CIFAR-100. Our method is marked with \blacklozenge .

Method	CIFAR-10 (Clean Acc.)		CIFAR-10 (Robust Acc.)		CIFAR-100 (Clean Acc.)		CIFAR-100 (Robust Acc.)	
	Overall	Worst	Overall	Worst	Overall	Worst	Overall	Worst
AT	83.5	67.3	47.9	19.6	53.7	21.0	19.7	1.6
FAIR-TAT \blacklozenge	85.6	73.7	46.7	19.7	61.4	25.0	20.9	2.2
AT + EMA	84.1	67.2	49.4	21.3	57.7	24.0	23.4	1.4
FAIR-TAT + EMA \blacklozenge	85.8	72.4	47.9	21.7	61.3	27.0	23.2	2.6
AT + CFA	83.8	68.1	50.1	22.8	59.3	24.8	24.7	2.4
FAIR-TAT+CFA \blacklozenge	84.8	72.0	48.8	24.6	62.8	27.6	25.0	2.9
TRADES	81.9	62.9	53.1	24.3	56.5	13.0	28.2	2.0
FAIR-TAT + TRADES \blacklozenge	81.7	68.1	51.9	26.2	59.4	24.0	29.4	8.9
TRADES + EMA	82.6	65.3	53.7	24.6	58.0	17.9	16.0	3.0
FAIR-TAT + TRADES + EMA \blacklozenge	81.9	68.9	53.3	26.8	59.4	26.0	27.7	9.1
TRADES + CFA	82.3	65.4	53.7	25.2	60.2	18.6	29.4	3.2
FAIR-TAT + TRADES + CFA \blacklozenge	82.1	68.6	53.1	26.9	59.7	26.4	29.6	8.6
FAT	84.9	68.5	48.1	18.7	54.0	16.0	19.3	1.8
FAIR-TAT + FAT \blacklozenge	86.6	77.2	46.1	22.7	60.9	24.0	20.9	2.7
FAT + EMA	86.3	69.2	48.5	18.8	56.9	22.0	23.6	1.6
FAIR-TAT + FAT + EMA \blacklozenge	86.8	73.5	47.5	21.3	61.6	23.2	21.9	2.8
FAT + CFA	85.0	71.0	51.0	23.4	59.8	24.9	24.1	2.0
FAIR-TAT + FAT + CFA \blacklozenge	86.1	74.9	48.2	24.4	60.0	25.3	23.1	2.3
FRL	82.8	71.4	45.7	24.4	52.0	23.0	22.7	1.2
FRL + EMA	83.6	69.5	46.3	24.8	54.0	24.2	22.5	1.0
BAT + TRADES	86.5	73.4	49.8	22.1	60.9	25.9	24.4	1.8
WAT + TRADES	81.2	65.9	47.1	26.3	57.3	21.2	25.6	2.9
Clean Training	94.0	79.4	13.7	2.7	76.7	39.2	4.1	0

ples. Additionally, we evaluate the overall robustness and fairness of our method concerning common corruptions.

Data. Our experiments utilize the benchmark datasets CIFAR-10 and CIFAR-100 [23]. We assess the performance of our approach under the following architectures: PRN-18, and XCiT-S12. These architectures are widely recognized for assessing robustness [8]. In the main paper, we present results based on CIFAR-10 and CIFAR100 using the PRN-18 architecture. Furthermore, we extend our evaluation to assess the model’s robustness and fairness against common corruptions, utilizing the CIFAR-10C dataset [19]. Notably, this study marks the first evaluation of common corruptions within the context of model fairness towards classification.

Metrics and Evaluations. To demonstrate the robustness and fairness of the model, we report both the overall average clean and robust accuracy, as well as the clean and robust accuracy of the hard class in the form of worst class accuracy.

Baselines. For our primary baseline, we select Vanilla Adversarial Training (AT) [30], TRADES [49] and FAT [51]. Given that the primary objective of our model is fairness, we compare our approach to FRL [45], which is the first adversarial training algorithm specifically designed to improve class-wise robustness and fairness. Additionally, we include CFA [44], a recent instance-based adaptive adversarial training approach that dynamically customizes perturbation mar-

gins during adversarial training. We also adapt CFA [44], TRADES [49] and FAT [51] according to our approach and compare the results with their respective baseline versions. We include recent fairness baselines such as BAT [38] and WAT [26]. Furthermore, we apply weight averaging techniques, such as Fairness Aware Weight Average (FAWA) [44] as in CFA and Exponential Moving Average (EMA) as well for our approach. Note that we have reported the results from the last checkpoint evaluations for all the methods therefore the results varies when reported with best checkpoint version. Also, The worst class accuracy for CIFAR-100 is defined as the average performance of the worst 10% of classes following the baseline works.

Training Settings. We train PRN-18 network in accordance with the settings outlined in [37]. Both networks undergo training using stochastic gradient descent (SGD) with a momentum of 0.9 and weight decay set at 5×10^{-4} . The learning rate scheduler is initially configured at 0.1. Learning rate adjustments are made by dividing the rate by 10 after 50% and 75% of the epochs during training. Adversarial training is conducted with a default perturbation margin of $\epsilon = 8/255$ and a robustness regularization parameter of $\beta = 2$. The parameter λ_1 is set to 0.5, which is used for scaling the customized perturbation of the classes. To enhance the initialization process, the weight averaging procedure

Table 2. Comparison of AutoAttack and Squares evaluations of FAIR-TAT framework with other methods focused on model fairness using PRN-18 on CIFAR-10 and CIFAR-100. Our method is marked with \blacklozenge .

Method	CIFAR-10 (AA)		CIFAR-10 (Squares)		CIFAR-100 (AA)	
	Overall	Worst Class	Overall	Worst Class	Overall	Worst Class
AT	45.7	15.4	51.3	19.4	17.8	1.3
FAIR-TAT \blacklozenge	45.0	18.7	51.8	26.7	19.5	1.9
AT + EMA	45.6	15.4	51.7	19.7	21.6	1.2
FAIR-TAT + EMA \blacklozenge	45.0	18.8	51.8	26.7	21.3	2.0
AT + CFA	47.4	19.3	51.7	19.9	21.9	1.8
FAIR-TAT+CFA \blacklozenge	47.0	24.1	52.6	31.8	22.3	2.1
TRADES	49.8	18.7	54.0	22.4	25.9	1.7
FAIR-TAT + TRADES \blacklozenge	48.6	20.8	52.6	26.1	26.7	4.8
TRADES + EMA	49.8	18.6	54.0	22.4	14.9	2.1
FAIR-TAT + TRADES + EMA \blacklozenge	48.6	21.1	52.6	26.1	24.9	4.0
TRADES + CFA	50.3	21.3	54.2	22.6	26.2	2.9
FAIR-TAT + TRADES + CFA \blacklozenge	49.8	24.0	53.0	25.8	24.7	4.7
FAT	44.2	16.0	50.8	20.8	17.6	1.3
FAIR-TAT + FAT \blacklozenge	44.6	17.9	51.5	26.3	18.4	1.8
FAT + EMA	44.2	15.9	50.8	20.8	20.3	1.3
FAIR-TAT + FAT+ EMA \blacklozenge	44.6	18.1	51.5	26.3	18.6	2.0
FAT + CFA	49.4	22.6	53.4	24.1	21.3	1.7
FAIR-TAT + FAT + CFA \blacklozenge	44.3	22.1	49.0	24.4	20.2	1.5
FRL	44.0	23.2	49.7	24.6	20.1	1.2
FRL + EMA	44.2	23.9	50.8	24.9	19.9	1.0
BAT + TRADES	44.1	18.7	51.5	25.0	21.0	1.0
WAT + TRADES	45.9	<u>24.3</u>	51.7	26.2	22.6	2.1

begins after the 50th epoch. For the FAWA weight averaging method, a validation set is constructed by utilizing 2% of samples from each class, while the remaining 98% of train set is employed for training. This approach ensures that FAWA does not impose a substantial computational burden. The fairness threshold for FAWA is established at 0.2. All the other baseline versions are trained using their respective training configurations. We conducted our experiments on internal cluster with 8 RTX A1000 GPUs.

Performance of Fairness and Robustness on PGD. We compare our approach, FAIR-TAT, with other methods aimed at improving robust model fairness. All baseline methods mentioned in Section 6, including ours, are trained using PGD. Table 1 presents the clean and robust accuracy results for the CIFAR-10 and CIFAR-100 dataset using the PRN-18 architecture, with robustness evaluated based on PGD attacks. In Table 1, "AT" refers to conventional adversarial training, while "FAIR-TAT" is our comparative adversarial approach. We also show adaptability of our approach to other methods for comparison, such as comparing TRADES with FAIR-TAT + TRADES. To be fair, each approach customizes perturbation margins for individual class during training. Results without customized perturbation margins are provided in the supplementary material (see section A.3).

From Table 1, we observe that FAIR-TAT significantly improves clean accuracy, both overall and for the worst class, in both CIFAR-10 and CIFAR-100 datasets. Additionally,

it enhances worst-class robust accuracy while maintaining nearly the same overall robustness, with only marginal reductions in case of CIFAR-10, and improves overall robustness in CIFAR-100. This improvement in CIFAR-100 can be attributed to the dataset’s diversity and larger number of classes, which enhances the ability to differentiate features more effectively. Overall, these results suggest that FAIR-TAT offers a better trade-off between robustness and fairness.

We also compare FAIR-TAT with weight averaging techniques such as EMA and FAWA, as used in CFA [44]. Our method maintains fairness after applying these techniques. Specifically, for the larger number of classes in CIFAR-100, FAIR-TAT with TRADES setting outperforms other baselines in both fairness and robustness, demonstrating that our method scales well with an increasing number of classes. Overall, our FAIR-TAT approach outperforms state-of-the-art methods in terms of clean accuracy and fairness when evaluated against the adversaries it was trained on, though with a minor trade-off in robust accuracy in certain cases.

Transferability to Other Adversaries. To evaluate the transferability of our approach, we extend our evaluation to adversaries that were not initially trained on. We use AutoAttack [9], a widely used ensemble adversarial attack and Squares attack [4], a black-box attack to evaluate the adversarial robustness to access our approach compared to baselines. We utilize the same baselines as in Table 1. Ta-

Table 3. Minimum Class accuracies of FAIR-TAT method on common corruptions. T: TRADES, F: FAT

Corruption Type (C)	FAIRTAT	FAIRTAT (EMA)	FAIRTAT (CFA)	FAIRTAT (T)	FAIRTAT (T+EMA)	FAIRTAT (T+CFA)	FAIRTAT (F)	FAIRTAT (F+EMA)	FAIRTAT (F+CFA)
gaussian_noise	0.63	0.61	0.58	0.60	0.60	0.58	0.58	0.63	<u>0.66</u>
shot_noise	0.66	0.64	0.60	0.61	0.61	0.61	0.61	0.66	<u>0.68</u>
speckle_noise	0.66	0.65	0.62	0.61	0.61	0.61	0.61	0.66	<u>0.67</u>
impulse_noise	0.51	0.49	0.52	0.39	0.46	0.66	0.66	<u>0.68</u>	0.64
defocus_blur	0.68	0.67	0.66	0.70	0.68	<u>0.71</u>	<u>0.71</u>	0.68	0.69
gaussian_blur	0.65	0.67	0.65	0.64	0.62	<u>0.68</u>	0.62	0.62	0.37
motion_blur	0.59	0.62	0.58	0.62	0.61	0.64	0.61	<u>0.67</u>	0.45
zoom_blur	0.67	<u>0.70</u>	0.64	0.65	0.63	0.64	0.64	0.64	0.40
snow	0.63	0.63	0.63	0.58	0.55	0.55	<u>0.64</u>	<u>0.64</u>	0.63
fog	0.45	0.39	0.34	<u>0.47</u>	0.20	0.17	0.45	0.32	0.15
brightness	0.63	<u>0.72</u>	0.65	0.65	0.63	0.66	<u>0.72</u>	0.67	<u>0.72</u>
contrast	0.22	0.20	0.17	0.22	0.20	0.18	<u>0.24</u>	0.18	0.17
elastic_transform	0.66	<u>0.67</u>	0.61	0.62	0.64	0.64	0.65	0.64	0.65
pixelate	0.66	0.68	0.63	0.64	0.66	0.63	0.68	0.68	<u>0.73</u>
jpeg_compression	0.66	<u>0.71</u>	0.63	0.64	0.63	0.67	0.68	0.68	<u>0.71</u>
spatter	<u>0.64</u>	<u>0.64</u>	0.60	0.62	0.62	0.62	0.62	0.63	0.63
saturate	<u>0.67</u>	0.66	0.61	0.63	0.63	0.66	0.64	0.66	0.66
frost	0.60	0.57	0.54	0.58	0.54	0.58	0.62	0.62	<u>0.63</u>

ble 2 showcases the evaluations of AutoAttack and Squares on CIFAR-10 and AutoAttack on CIFAR-100 using PRN-18. Standard AutoAttack itself is an ensemble that contains Squares attack, however, we present the Squares evaluation separately to show the performance of our method when evaluated on a black-box attack alone. From the evaluations, it is clear that our approach, FAIR-TAT with respective weight averaging schemes, provides better transferability of fairness by enhancing the worst class accuracy while maintaining the marginal overall model robustness when assessed against different adversaries outside of the training domain.

Table 4. Minimum Class accuracies of baseline methods on common corruptions. T: TRADES, F: FAT

Corruption Type (C)	AT	AT (EMA)	AT (CFA)	AT (T)	AT (T+EMA)	AT (T+CFA)	AT (F)	AT (F+EMA)	AT (F+CFA)
gaussian_noise	0.57	0.58	0.54	0.52	0.55	0.54	<u>0.66</u>	0.63	0.44
shot_noise	0.58	0.60	0.56	0.53	0.57	0.56	<u>0.68</u>	0.66	0.45
speckle_noise	0.58	0.60	0.56	0.54	0.57	0.56	<u>0.67</u>	0.66	0.44
impulse_noise	0.52	0.53	0.49	0.46	0.51	0.52	0.64	0.66	<u>0.68</u>
defocus_blur	0.64	0.66	0.70	0.62	0.63	0.64	<u>0.71</u>	0.68	0.69
gaussian_blur	0.63	0.66	<u>0.68</u>	0.59	0.64	0.59	0.62	0.62	0.40
motion_blur	0.62	0.62	<u>0.63</u>	0.57	0.62	0.58	0.62	0.61	0.45
zoom_blur	0.64	0.67	<u>0.68</u>	0.59	0.65	0.59	0.64	0.68	0.40
snow	0.53	0.59	0.55	0.56	0.58	<u>0.64</u>	0.63	0.63	0.63
fog	0.46	0.43	0.43	0.42	<u>0.47</u>	0.34	0.45	0.32	0.15
brightness	0.62	0.67	<u>0.72</u>	0.63	0.65	<u>0.72</u>	0.67	0.71	<u>0.72</u>
contrast	<u>0.22</u>	0.17	0.17	0.24	0.18	0.18	0.18	0.17	<u>0.22</u>
elastic_transform	<u>0.66</u>	0.61	0.64	0.65	0.62	0.64	0.64	0.65	0.65
pixelate	0.66	0.63	0.68	0.68	0.68	0.63	0.68	<u>0.73</u>	0.64
jpeg_compression	0.66	0.63	0.67	0.68	0.68	0.67	0.68	<u>0.71</u>	0.64
spatter	<u>0.64</u>	0.60	0.62	0.62	0.62	0.62	0.63	0.63	0.62
saturate	<u>0.67</u>	0.61	0.66	0.66	0.63	0.66	0.66	0.66	0.66
frost	0.60	0.54	0.58	0.62	0.62	0.62	<u>0.63</u>	0.62	0.62

Fairness and Robustness on Common Corruptions. In this work, we evaluate the robustness and fairness of models trained using our approach on common corruptions compared to baseline methods. We use the CIFAR-10C dataset, which contains 18 types of corruptions applied to CIFAR-10 images. Table 3 and Table 4 present the worst class accuracies for our method and the baselines, respectively. Table 3 shows the minimum class accuracies of FAIR-TAT adapted to other approaches, while Table 4 lists the minimum class accuracies of the baselines. The best results on common corruptions of baselines and our methods are underlined in

the tables. Our approach FAIR-TAT, when adapted to baselines, outperforms them in terms of worst class accuracies and overall model performance. The overall accuracies on common corruptions are provided in the supplementary (see section B.3). Overall, our findings suggest that FAIR-TAT improves both robustness and fairness compared to existing methods, with rare exceptions. Although FAT performs comparably equally well on common corruptions, our approach achieves a better balance in performance considering fairness on corruptions as well on adversarial samples.

Discussion of Limitations. FAIR-TAT enhances worst-class accuracy, but there’s still significant room to improve class-wise balance. We see our method as a step toward making models both robust and fair. While FAIR-TAT offers slightly better trade-offs between clean and robust accuracy than traditional AT, clean accuracy still declines. Technically, further exploration is needed; for example, we currently sample targets based on class-wise misclassification using C_{FPS} . Although focusing on sample-wise misclassifications instead of class-wise misclassifications would be more beneficial, it is costly, and we plan to address this in future work.

Broader Impact. Ensuring both robustness and fairness in neural models is important for their deployment in safety-critical applications. Our work focuses in the direction of training a robust and fair classifier. Robust and fair models maintain reliability under diverse conditions, reducing the risk of unexpected failures. Fair models ensure equitable treatment across all demographics, preventing bias-related harm. Although improvements in robustness of classifiers are desirable, they can also lead to a false sense of security.

7. Conclusion

In conclusion, our empirical analysis highlights the need to focus more on specific classes which are prone to be misclassified often than others. Building upon these findings, we introduce the Fair Targeted Adversarial Training (FAIR-TAT) framework which samples these most confused classes more often than others and learn to defend these class-specific perturbations by generating targeted adversaries. This allows to learn the class-to-class bias more effectively. FAIR-TAT demonstrates improved fairness against adversarial attacks compared to state-of-the-art baseline models not only on the attacks it is trained on but also on other type of adversaries. Simultaneously, FAIR-TAT enhances accuracy on clean samples and common corruptions, outperforming existing baselines showcasing the overall effectiveness and balance of our approach against diverse threats.

Acknowledgements

Steffen Jung and Margret Keuper acknowledge funding by the DFG Research Unit 5336 - Learning to Sense.

References

- [1] Shashank Agnihotri, Kanchana Vaishnavi Gandikota, Julia Grabinski, Paramanand Chandramouli, and Margret Keuper. On the unreasonable vulnerability of transformers for image restoration-and an easy fix. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3707–3717, 2023. 2
- [2] Shashank Agnihotri, Julia Grabinski, and Margret Keuper. Improving feature stability during upsampling – spectral artifacts and the importance of spatial context. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol, editors, *Computer Vision – ECCV 2024*, pages 357–376, Cham, 2025. Springer Nature Switzerland. 2
- [3] Shashank Agnihotri, Steffen Jung, and Margret Keuper. Cospgd: a unified white-box adversarial attack for pixel-wise prediction tasks. *arXiv preprint arXiv:2302.02213*, 2023. 1, 2
- [4] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: A query-efficient black-box adversarial attack via random search. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, Lecture Notes in Computer Science, pages 484–501. Springer International Publishing, 2020. 3, 7
- [5] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 274–283. PMLR, 10–15 Jul 2018. 2
- [6] Philipp Benz, Chaoning Zhang, Adil Karjauv, and In So Kweon. Robustness may be at odds with fairness: An empirical study on class-wise accuracy. In *Preregister@NeurIPS*, 2020. 1, 3
- [7] Nicholas Carlini and David Wagner. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pages 3–14. ACM, 2017. 3
- [8] Francesco Croce, Maksym Andriushchenko, Vikash Sehwal, Edoardo DeBenedetti, Nicolas Flammarion, Mung Chiang, Prateek Mittal, and Matthias Hein. RobustBench: a standardized adversarial robustness benchmark. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2021. 2, 4, 6
- [9] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International conference on machine learning*, pages 2206–2216. PMLR, 2020. 7
- [10] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *Proceedings of the 37th International Conference on Machine Learning*, pages 2206–2216. PMLR, 2020-11-21. ISSN: 2640-3498. 3
- [11] Francesco Croce and Matthias Hein. Mind the box: l_1 -apgd for sparse adversarial attacks on image classifiers. In *International Conference on Machine Learning*, pages 2201–2211. PMLR, 2021. 1, 2
- [12] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 1
- [13] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples, 2015. 1, 2, 3
- [14] Sven Gowal, Sylvestre-Alvise Rebuffi, Olivia Wiles, Florian Stimberg, Dan Andrei Calian, and Timothy Mann. Improving robustness using generated data. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. 2
- [15] Julia Grabinski, Paul Gavrikov, Janis Keuper, and Margret Keuper. Robust models are less over-confident. *Advances in Neural Information Processing Systems*, 35:39059–39075, 2022. 2, 4
- [16] Julia Grabinski, Janis Keuper, and Margret Keuper. Aliasing and adversarial robust generalization of cnns. *Machine Learning*, 111(11):3925–3951, 2022. 2
- [17] Julia Grabinski, Janis Keuper, and Margret Keuper. As large as it gets-studying infinitely large convolutions via neural implicit frequency filters. *Transactions on Machine Learning Research*, 2024:1–42, 2024. 2
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. ISSN: 1063-6919. 1
- [19] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019. 1, 2, 6
- [20] Steffen Jung, Jovita Lukasik, and Margret Keuper. Neural architecture design and robustness: A dataset. In *ICLR*, 2023. 2
- [21] Oğuzhan Fatih Kar, Teresa Yeo, Andrei Atanov, and Amir Zamir. 3D Common Corruptions and Data Augmentation. In *3D Common Corruptions and Data Augmentation*, pages 18963–18974, 2022. 1
- [22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 1
- [23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>, 2009. 6
- [24] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale, 2017. 1, 2
- [25] Hyungyu Lee, Saehyung Lee, Hyemi Jang, Junsung Park, Ho Bae, and Sungroh Yoon. Dafa: Distance-aware fair adversarial training. *arXiv preprint arXiv:2401.12532*, 2024. 3
- [26] Boqi Li and Weiwei Liu. Wat: improve the worst-class robustness in adversarial training. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 14982–14990, 2023. 1, 3, 6

- [27] Jovita Lukasik, Paul Gavrikov, Janis Keuper, and Margret Keuper. Improving native cnn robustness with filter frequency regularization. *Transactions on Machine Learning Research*, 2023:1–36, 2023. [2](#)
- [28] Xingjun Ma, Yuhao Niu, Lin Gu, Yisen Wang, Yitian Zhao, James Bailey, and Feng Lu. Understanding adversarial attacks on deep learning based medical image analysis systems. *Pattern Recognition*, 110:107332, 2021. [1](#)
- [29] Xinsong Ma, Zekai Wang, and Weiwei Liu. On the trade-off between robustness and fairness. *Advances in Neural Information Processing Systems*, 35:26230–26241, 2022. [1](#)
- [30] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. [1](#), [2](#), [3](#), [6](#)
- [31] Tejaswini Medi, Julia Grabinski, and Margret Keuper. Towards class-wise robustness analysis, 2024. [2](#)
- [32] Claudio Michaelis, Benjamin Mitzkus, Robert Geirhos, Evgenia Rusak, Oliver Bringmann, Alexander S. Ecker, Matthias Bethge, and Wieland Brendel. Benchmarking robustness in object detection: Autonomous driving when winter is coming. <https://openreview.net/forum?id=ryljMpNtwr>, 2020. [1](#)
- [33] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. DeepFool: A simple and accurate method to fool deep neural networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2574–2582. IEEE, 2016. [3](#)
- [34] Tianyu Pang, Xiao Yang, Yinpeng Dong, Kun Xu, Jun Zhu, and Hang Su. Boosting adversarial training with hypersphere embedding, 2020. [2](#)
- [35] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John C. Duchi, and Percy Liang. Adversarial training can hurt generalization, 2019. [2](#)
- [36] Sylvestre-Alvise Rebuffi, Sven Gowal, Dan Andrei Calian, Florian Stimberg, Olivia Wiles, and Timothy Mann. Data augmentation can improve robustness. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. [2](#)
- [37] Leslie Rice, Eric Wong, and Zico Kolter. Overfitting in adversarially robust deep learning. In *International Conference on Machine Learning*, pages 8093–8104. PMLR, 2020. [2](#), [6](#)
- [38] Chunyu Sun, Chenye Xu, Chengyuan Yao, Siyuan Liang, Yichao Wu, Ding Liang, Xianglong Liu, and Aishan Liu. Improving robust fairness via balance adversarial training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 15161–15169, 2023. [1](#), [2](#), [6](#)
- [39] Qi Tian, Kun Kuang, Kelu Jiang, Fei Wu, and Yisen Wang. Analysis and applications of class-wise robustness in adversarial training. In Feida Zhu, Beng Chin Ooi, and Chunyan Miao, editors, *KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, Singapore, August 14-18, 2021*, pages 1561–1570. ACM, 2021. [2](#), [3](#)
- [40] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. *arXiv preprint arXiv:1805.12152*, 2018. [1](#)
- [41] Hongjun Wang and Yisen Wang. Generalist: Decoupling natural and robust generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20554–20563, 2023. [1](#)
- [42] Yisen Wang, Xingjun Ma, James Bailey, Jinfeng Yi, Bowen Zhou, and Quanquan Gu. On the convergence and robustness of adversarial training. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6586–6595. PMLR, 09–15 Jun 2019. [1](#), [2](#)
- [43] Zekai Wang, Tianyu Pang, Chao Du, Min Lin, Weiwei Liu, and Shuicheng Yan. Better diffusion models further improve adversarial training. *arXiv preprint arXiv:2302.04638*, 2023. [1](#), [2](#)
- [44] Zeming Wei, Yifei Wang, Yiwen Guo, and Yisen Wang. Cfa: Class-wise calibrated fair adversarial training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8193–8201, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [45] Han Xu, Xiaorui Liu, Yaxin Li, Anil Jain, and Jiliang Tang. To be robust or to be fair: Towards fairness in adversarial training. In *International conference on machine learning*, pages 11492–11501. PMLR, 2021. [1](#), [2](#), [6](#)
- [46] Dongyoon Yang, Insung Kong, and Yongdai Kim. Improving adversarial robustness by putting more regularizations on less robust samples. In *International Conference on Machine Learning*, pages 39331–39348. PMLR, 2023. [1](#)
- [47] Xiaoyong Yuan, Pan He, Qile Zhu, and Xiaolin Li. Adversarial examples: Attacks and defenses for deep learning. *IEEE Transactions on Neural Networks and Learning Systems*, 30:2805–2824, 2017. [2](#)
- [48] Xinli Yue, Mou Ningping, Qian Wang, and Lingchen Zhao. Revisiting adversarial robustness distillation from the perspective of robust fairness. *Advances in Neural Information Processing Systems*, 36, 2024. [1](#)
- [49] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *International conference on machine learning*, pages 7472–7482. PMLR, 2019. [1](#), [6](#)
- [50] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 7472–7482. PMLR, 09–15 Jun 2019. [2](#)
- [51] Jingfeng Zhang, Xilie Xu, Bo Han, Gang Niu, Lizhen Cui, Masashi Sugiyama, and Mohan Kankanhalli. Attacks which do not kill training make adversarial learning stronger. In *International conference on machine learning*, pages 11278–11287. PMLR, 2020. [6](#)
- [52] Shiji Zhao, Xizhe Wang, and Xingxing Wei. Improving adversarial robust fairness via anti-bias soft label distillation. *arXiv preprint arXiv:2312.05508*, 2023. [1](#)