

Recent Advances in Adversarial Training for Adversarial Robustness

Tao Bai^{1*}, Jinqi Luo¹, Jun Zhao¹, Bihan Wen¹ and Qian Wang²

¹Nanyang Technological University, Singapore

²Wuhan University, China

{bait0002, luoj0021, junzhao, bihan.wen}@ntu.edu.sg, qianwang@whu.edu.cn

Abstract

Adversarial training is one of the most effective approaches for deep learning models to defend against adversarial examples. Unlike other defense strategies, adversarial training aims to enhance the robustness of models intrinsically. During the past few years, adversarial training has been studied and discussed from various aspects, which deserves a comprehensive review. For the first time in this survey, we systematically review the recent progress on adversarial training for adversarial robustness with a novel taxonomy. Then we discuss the generalization problems in adversarial training from three perspectives and highlight the challenges which are not fully tackled. Finally, we present potential future directions.

1 Introduction

The adversarial vulnerability of deep neural networks has attracted significant attention in recent years. With slight but carefully-crafted perturbations, the perturbed natural images, namely adversarial examples [Szegedy *et al.*, 2014], can mislead state-of-the-art (SOTA) classifiers to make erroneous predictions. Besides classification, adversarial examples appear in various tasks like semantic segmentation, object detection, and super-resolution (see the summary in [Yuan *et al.*, 2019]). The existence of adversarial examples raises concerns from the public and motivates the proposals of defenses [Goodfellow *et al.*, 2015; Huang *et al.*, 2015; Madry *et al.*, 2018]. Naturally, the defenses also stimulate the development of stronger attacks, seeming like an arms race.

Among various existing defense strategies, Adversarial Training (AT) [Goodfellow *et al.*, 2015; Madry *et al.*, 2018] proves to be the most effective against adversarial attacks [Pang *et al.*, 2021; Maini *et al.*, 2020; Schott *et al.*, 2019], receiving considerable attention from the research community. The idea of adversarial training is straightforward: it augments training data with adversarial examples in each training loop. Thus adversarially trained models behave more normally when facing adversarial examples than standardly trained models. Mathematically, adversarial training is

formulated as a min-max problem, searching for the best solution to the worst-case optimum. The main challenge of adversarial training is to solve the inner maximization problem, which researchers are actively working on. The last few years have witnessed tremendous efforts made by the research community. The recent advances have resulted in a variety of new techniques in the literature, which no doubt deserves a comprehensive review.

To our best knowledge, surveys focusing on adversarial training do not exist so far. Existing surveys are aimed at providing a complete review of adversarial machine learning and summarize all the existing adversarial attacks and defense methods [Yuan *et al.*, 2019; Silva and Najafirad, 2020; Sadeghi *et al.*, 2020; Tabassi *et al.*, 2019; Smith, 2019], which inevitably omit important details of subfields. In contrast, our goal in this paper is to give a detailed overview of adversarial training in image classification, which has attracted extensive attention of researchers. We believe that this survey can provide up-to-date findings and developments happening on adversarial training. Notably, we carefully review and analyze adversarial training with a novel taxonomy, uniquely discuss the poor generalization ability from different perspectives, and present future research directions.

2 Preliminaries

Adversarial Attacks. Adversarial attacks refer to finding adversarial examples for well-trained models. In this paper, we consider only the situation where the training/test data are initially from the same distribution. Taking classification as an example, we use $f(x; \theta) : \mathbb{R}^{h \times w \times c} \rightarrow \{1 \dots k\}$ to denote an image classifier that maps an input image x to a discrete label set C with k classes, in which θ indicates the parameters of f , and h, w, c represent image height, width and channel, respectively. Given the perturbation budget ϵ , the adversary tries to find a perturbation $\delta \in \mathbb{R}^{h \times w \times c}$ to maximize the loss function, e.g., cross-entropy loss \mathcal{L}_{ce} , so that $f(x + \delta) \neq f(x)$. Therefore, δ is estimated as

$$\delta^* := \arg \max_{|\delta|_p \leq \epsilon} \mathcal{L}_{ce}(\theta, x + \delta, y), \quad (1)$$

where y is the label of x , and p can be 0, 1, 2 and ∞ . In most cases, ϵ is small so that the perturbations are imperceptible to human eyes. Note that we only consider the l_p -based attacks

*Contact Author

for classification in this paper. The adversarial counterpart x' of x is expressed as

$$x' := x + \delta^*. \quad (2)$$

There are some common attacks following this formulation, such as Fast Gradient Sign Method (FGSM) [Goodfellow *et al.*, 2015], iterative FGSM [Kurakin *et al.*, 2016], and Projected Gradient Descent (PGD) attack [Madry *et al.*, 2018].

Adversarial Robustness. In general, adversarial robustness is the model’s performance on test data with adversarial attacks, *i.e.*, adversarial test accuracy of classification. We expect to know the model’s performance in the worst-case scenario, so the adversary should be strong enough and launch attacks in white-box settings, *i.e.*, the adversary has full knowledge of the model, such as architectures, parameters, training data, etc. In practice, PGD attack [Madry *et al.*, 2018] is commonly employed for evaluation because of its strong attack ability in white-box settings.

3 Adversarial Training for Adversarial Robustness

Currently, adversarial training is widely accepted as the most effective method in practice to improve the adversarial robustness of deep learning models [Athalye and Carlini, 2018]. However, it is still a long way to go for adversarial training to handle adversarial attacks perfectly. Prevailing adversarial training methods [Madry *et al.*, 2018] can produce a robust model with worst-case accuracy of around 90% on MNIST. For slightly more challenging datasets, *e.g.*, CIFAR-10, adversarial training only achieves around 45% and 40% on SVHN [Buckman *et al.*, 2018], which is far from satisfactory. Additionally, adversarial training leads to the degraded generalization ability of deep learning models. In this section, we review the development of adversarial training, summarize the recent advances with a novel taxonomy, and discuss the generalization problem of adversarial training.

3.1 The Origin of Adversarial Training

The initial idea of adversarial training is firstly brought to light by [Dalvi *et al.*, 2004; Lowd and Meek, 2005] in a game-theoretic setup. Szegedy *et al.* (2014), for the first time, adopted this idea to obtain robust neural networks, where neural networks are trained on a mixture of adversarial examples and clean examples. Goodfellow *et al.* (2015) went further and proposed FGSM to produce adversarial examples during training. Yet, their trained models remain vulnerable to iterative attacks [Tramèr *et al.*, 2018] as these approaches utilized a linear function to approximate the loss function, leading to sharp curvature near data points on the decision surface of the corresponding deep models. The existence of sharp curvature is also known as gradient masking [Papernot *et al.*, 2017].

Unlike the prior works in which models are trained on a mixture of clean data and adversarial data, a line of research trains models with adversarial data only. For the first time, Huang *et al.* (2015) defined a min-max problem that the training procedure is forced to minimize classification error against an adversary who perturbs the input and maximizes the classification error. They also pointed out that the key

to solving this min-max problem is finding strong adversarial examples. Shiham *et al.* (2018) considered this min-max problem from a robust optimization perspective and proposed the framework of adversarial training. The formulation is illustrated below:

$$\min_{\theta} \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\max_{\delta \in B(x,\varepsilon)} \mathcal{L}_{ce}(\theta, x + \delta, y) \right], \quad (3)$$

where $(x, y) \sim \mathcal{D}$ represents training data sampled from distribution \mathcal{D} and $B(x, \varepsilon)$ is the allowed perturbation set, expressed as $B(x, \varepsilon) := \{x + \delta \in \mathbb{R}^{h \times w \times c} \mid \|\delta\|_p \leq \varepsilon\}$. Madry *et al.* (2018) gave a reasonable interpretation of this formulation: the inner maximization problem is finding the worst-case samples for the given model, and the outer minimization problem is to train a model robust to adversarial examples.

With such connection, Madry *et al.* (2018) employed a multi-step gradient based attack known as PGD attack for solving the inner problem as follows:

$$x^{t+1} := \text{Proj}_{x+B(x,\varepsilon)} (x^t + \alpha \text{sign}(\nabla_{x^t} \mathcal{L}_{ce}(\theta, x^t, y))), \quad (4)$$

where t is the current step and α is the step size. Further, they investigated the inner maximization problem from the landscape of adversarial examples and gave both theoretical and empirical proofs of local maxima’s tractability with PGD. Through extensive experiments, their approach (PGD-AT) significantly increased the adversarial robustness of deep learning models against a wide range of attacks, which is a milestone of adversarial training methods. As most derivative works followed their designs and settings, PGD-AT became a critical benchmark and is regarded as the standard way to do adversarial training in practice.

3.2 Taxonomy of Adversarial Training

In this subsection, we review the recent advances of adversarial training in last few years, categorized by different understandings on adversarial training. A summary of selected adversarial training methods is provided in Table 1.

Adversarial Regularization

The idea of adversarial regularization first appears in [Goodfellow *et al.*, 2015]. Besides cross-entropy loss, they added a regularization term in the objective function, which is based on FGSM and expressed as $\mathcal{L}(\theta, x + \epsilon \text{sign}(\nabla_x \mathcal{L}(\theta, x, y)))$. Kurakin *et al.* (2017) extended this FGSM-based regularization term by controlling the ratio of adversarial examples in batches so that it can scale up to ImageNet. The effectiveness of their method is validated on single-step attacks as they assume the linearity of neural networks is attributed to the existence of adversarial examples [Goodfellow *et al.*, 2015]. However, Qin *et al.* (2019) calculated the absolute difference between the adversarial loss and its first-order Taylor expansion, concluding that more robust models usually have smaller values of local linearity. Correspondingly, they replaced the FGSM-based regularization with a Local Linearity Regularization for adversarial robustness.

Distinct from previous methods, [Zhang *et al.*, 2019b] decomposed the robust error \mathcal{R}_{rob} as the sum of natural error

Taxonomy	Publication	Model Architecture	Attack	ϵ	Dataset	Accuracy
Adversarial Regularization	[Qin <i>et al.</i> , 2019]	ResNet-152	PGD ₅₀	4/255	ImageNet	47.00%
	[Zhang <i>et al.</i> , 2019b]	Wide ResNet	CW ₁₀	0.031/1	CIFAR-10	84.03%
	[Wang <i>et al.</i> , 2020]	ResNet-18	PGD ₂₀	8/255	CIFAR-10	55.45%
	[Kannan <i>et al.</i> , 2018]	InceptionV3	PGD ₁₀	16/255	ImageNet	27.90%
	[Mao <i>et al.</i> , 2019]	Wide ResNet	PGD ₂₀	8/255	CIFAR-10	50.03%
Curriculum	[Zhang <i>et al.</i> , 2020]	Wide ResNet	PGD ₂₀	16/255	CIFAR-10	49.86%
	[Cai <i>et al.</i> , 2018]	DenseNet-161	PGD ₇	8/255	CIFAR-10	69.27%
	[Wang <i>et al.</i> , 2019]	8-Layer ConvNet	PGD ₂₀	8/255	CIFAR-10	42.40%
Ensemble	[Pang <i>et al.</i> , 2019]	Wide ResNet	PGD ₁₀	0.005	CIFAR-100	32.10%
	[Kariyappa and Qureshi, 2019]	ResNet-20	PGD ₃₀	0.09/1	CIFAR-10	46.30%
	[Yang <i>et al.</i> , 2020a]	ResNet-20	PGD ₂₀	0.01/1	CIFAR-10	52.40%
Adaptive ϵ	[Balaji <i>et al.</i> , 2019]	ResNet-152	PGD ₁₀₀₀	8/255	ImageNet	59.28%
	[Ding <i>et al.</i> , 2020]	Wide ResNet	PGD ₁₀₀	8/255	CIFAR-10	47.18%
	[Cheng <i>et al.</i> , 2020]	Wide ResNet	PGD ₂₀	8/255	CIFAR-10	73.38%
Semi-Unsupervised	[Alayrac <i>et al.</i> , 2019]	Wide ResNet	FGSM	8/255	CIFAR-10	62.18%
	[Carmon <i>et al.</i> , 2019]	Wide ResNet	PGD ₁₀	8/255	CIFAR-10	63.10%
	[Zhai <i>et al.</i> , 2019]	Customized ResNet	PGD ₇	8/255	CIFAR-10	42.48%
	[Hendrycks <i>et al.</i> , 2019]	Wide ResNet	PGD ₂₀	0.3/1	ImageNet	50.40%
Efficient	[Shafahi <i>et al.</i> , 2019]	Wide ResNet	PGD ₁₀₀	8/255	CIFAR-10	46.19%
	[Wong <i>et al.</i> , 2020]	ResNet-50	PGD ₄₀	2/255	ImageNet	43.43%
	[Andriushchenko and Flammarion, 2020]	ResNet-50	PGD ₅₀	2/255	ImageNet	41.40%
	[Kim <i>et al.</i> , 2021]	PreActResNet-18	FGSM	8/255	CIFAR-10	50.50%
	[S. and Babu, 2020]	Wide ResNet	PGD ₄₀	8/255	MNIST	88.51%
	[Song <i>et al.</i> , 2019]	Customized ConvNet	PGD ₂₀	4/255	CIFAR-10	58.10%
	[Vivek and Babu, 2020]	Wide ResNet	PGD ₁₀₀	0.3/1	MNIST	90.03%
	[Huang <i>et al.</i> , 2020]	Wide ResNet	PGD ₂₀	8/255	CIFAR-10	45.80%
[Zhang <i>et al.</i> , 2019a]	Wide ResNet	PGD ₂₀	8/255	CIFAR-10	47.98%	
Others	[Dong <i>et al.</i> , 2020]	Wide ResNet	PGD ₂₀	8/255	CIFAR-100	29.40%
	[Wang and Zhang, 2019]	Wide ResNet	CW ₂₀₀	4/255	CIFAR-10	60.30%
	[Zhang and Wang, 2019]	Wide ResNet	PGD ₂₀	8/255	CIFAR-100	47.20%
	[Pang <i>et al.</i> , 2020]	Wide ResNet	PGD ₅₀₀	8/255	CIFAR-10	60.75%
	[Lee <i>et al.</i> , 2020]	PreActResNet-18	PGD ₂₀	8/255	Tiny ImageNet	20.31%
Benchmark	[Madry <i>et al.</i> , 2018]	ResNet-50	PGD ₂₀	8/255	CIFAR-10	45.80%

Table 1: A summary of experimental results for various adversarial training methods. All the attacks are under l_∞ norm.

\mathcal{R}_{nat} and boundary error \mathcal{R}_{db} . Boundary error occurs when the distance between data and the decision boundary is sufficiently small (less than ϵ), which is also the reason for adversarial examples’ existence. So they proposed TRADES to minimize the \mathcal{R}_{db} by solving the following problem:

$$\min_f \mathbb{E} \{ \mathcal{L}_{ce}(f(x), y) + \max_{x' \in \mathbb{B}(x, \epsilon)} \mathcal{L}_{ce}(f(x), f(x')) / \lambda \}, \quad (5)$$

where λ is a coefficient, determining the strength of regularization. Such decomposition is proved to be effective, and TRADES outperforms PGD-AT on CIFAR-10 with error rates reduced by 10%. One problem of TRADES is that the regularization term is designed to push natural examples and their adversarial counterparts together, no matter natural data are classified correctly or not. [Wang *et al.*, 2020] investigated the influence of misclassified examples

and proposed Misclassification Aware adversarial Training (MART), which emphasizes on misclassified examples with weights $1 - \mathcal{P}_y(x, \theta)$, where $\mathcal{P}_y(x, \theta)$ is the probability of ground truth label y .

Due to the amplification of deep models, imperceptible noises could lead to substantial changes in feature space [Goodfellow *et al.*, 2015]. Some works analyze adversarial training from the perspective of representation. Kannan *et al.* (2018) proposed Adversarial Logit Pairing (ALP), encouraging logits for pairs of examples to be similar. But ALP initially is not useful due to the wrong formulation of adversarial training objectives [Engstrom *et al.*, 2018]. Further to enhance the alignment of representations of natural data and their adversarial counterparts, Mao *et al.* (2019) adopted the prevalent triplet loss for regularization, which uses adversar-

ial examples as anchors.

Adversarial regularization is an essential variant of adversarial training [Shaham *et al.*, 2018]. Compared to the original formulation of adversarial training, adversarial regularization is more flexible and requires a deep understanding of adversarial robustness. Also, the decomposition of robust error indeed paves the way for unlabeled data to enhance adversarial robustness.

Curriculum-based Adversarial Training

According to the formulation of adversarial training, the inner problem is always trying to find the worst-case samples. One natural question is: are those worst-case samples always suitable for adversarial training? Zhang *et al.* (2020) found that adversarial examples generated by strong attacks significantly cross over the decision boundary and are close to natural data. As PGD-AT only utilizes adversarial examples for training, this leads to overfitting of adversarial examples [Cai *et al.*, 2018].

For alleviating the overfitting, researchers adapt the idea of curriculum training to adversarial training. Cai *et al.* (2018) proposed Curriculum Adversarial Training (CAT), with an assumption that PGD with more steps generates stronger adversarial examples. Starting from a small number of steps, CAT gradually increases the iteration steps of PGD until the model achieves a high accuracy against the current attack. Different from CAT, Friendly Adversarial Training (FAT) [Zhang *et al.*, 2020] adapts early stopping when performing PGD attacks and returns adversarial data near the decision boundary for training. Both CAT and FAT adjust the attacks' strength in a practical way, where a quantitative criterion is missing. From the convergence point of view, Wang *et al.* (2019) designed First-Order Stationary Condition (FOSC) to estimate the convergence quality of the inner maximization problem. The closer the FOSC to zero, the stronger the attack.

Such curriculum-based methods help improve the generalization of clean data while preserving adversarial robustness. One possible reason for their success is weak attacks in early training stages are associated with generalization [Wang *et al.*, 2019]. In addition to relieving overfitting, curriculum-based methods reduce training time due to the varying iteration numbers of PGD for solving the inner maximization problem.

Ensemble Adversarial Training

Tramèr *et al.* (2018) firstly introduced ensemble learning into adversarial training, called Ensemble Adversarial Training (EAT), where training data is augmented with adversarial examples generated from different target models instead of a single model. The advantage of EAT is that it helps alleviate the sharp curvatures caused by the single-step attacks *e.g.*, FGSM. However, the interaction among different target models is neglected [Tramèr *et al.*, 2018]. Specifically, standardly trained target models may have similar predictions or representations [Dauphin *et al.*, 2014] and share the adversarial subspace [Tramèr *et al.*, 2017], which potentially hurts the performance of EAT.

For promoting the diversity among target models, several improvements are proposed, such as the adaptive diversity promoting regularizer [Pang *et al.*, 2019], forcing different

models to be diverse in non-maximal predictions; maximizing the cosine distances between target models' input gradients [Kariyappa and Qureshi, 2019] (input gradients refer to the gradients of the loss function *w.r.t.* the input); and maximizing the vulnerability diversity [Yang *et al.*, 2020a], which is defined as the sum of losses for two models with crafted images containing non-robust features [Ilyas *et al.*, 2019].

Intrinsically, such ensemble methods are useful for approximating the optimal value of the inner maximization problem in adversarial training. As proved in [Tramèr *et al.*, 2018], models trained with EAT have better generalization abilities regardless of the perturbation types. To conclude, adding the number and diversity of target models in training is a practical and useful way to approximate the space of adversarial examples, which is challenging to be described explicitly.

Adversarial Training with Adaptive ϵ

As shown in Equation (3), the parameters of attacks are predefined and fixed during training, such as ϵ . Some works [Balaji *et al.*, 2019; Ding *et al.*, 2020] argued individual data points might have different intrinsic robustness, *i.e.*, different distances to the classifier's decision boundary; however, adversarial training with fixed ϵ treats all data equally.

Considering the individual characteristic of adversarial robustness, researchers propose to do adversarial training at the instance level. Balaji *et al.* (2019) firstly presented Instance Adaptive Adversarial Training (IAAT), where ϵ is selected to be as large as possible, ensuring images within ϵ -ball of x are from the same class. This strategy helps IAAT relieve the trade-off between robustness and accuracy, though there is a slight drop in robustness. Unlike IAAT, another work called Margin Maximization Adversarial Training (MMA) [Ding *et al.*, 2020] directly maximizes the margin-distances between data points and the model's decision boundary, which is estimated by the adversarial perturbations with the least magnitudes. The manner of choosing ϵ in MMA is more reasonable as ϵ is sufficiently small, and such small ϵ in spatial domain hardly changes the classes of images substantially, especially for high-resolution images. The following work, Customized Adversarial Training (CAT) [Cheng *et al.*, 2020] further applies adaptive label uncertainty to prevent over-confident predictions based on adaptive ϵ .

Adversarial training with adaptive ϵ is a good exploration. However, empirical evidence shows many standard datasets are distributionally separated, *i.e.*, the distances inter classes are larger than ϵ used for attacks [Yang *et al.*, 2020b]. This reflects the limitation of current adversarial training methods on finding proper decision boundaries.

Adversarial Training with Semi/Unsupervised Learning

One key observation in supervised adversarial training methods [Madry *et al.*, 2018; Zhang *et al.*, 2019b] is adversarial accuracy in testing is much lower than in training. There is a large generalization gap in adversarial training (see Figure 1 in [Schmidt *et al.*, 2018]). The recent work [Schmidt *et al.*, 2018] studied this problem from the perspective of sample complexity. It is theoretically proved that adversarially robust training requires substantially larger datasets than standard training. However, high-quality datasets with labels are expensive to collect, which is of particular interest in practice.

Alternatively, several works appeared concurrently, exploring the possibility of training with additional unlabeled data.

Following the analysis of Gaussian models in [Schmidt *et al.*, 2018], a couple of works [Alayrac *et al.*, 2019; Carmon *et al.*, 2019; Zhai *et al.*, 2019] theoretically show that unlabeled data significantly reduces the sample complexity gap between standard training and adversarial training. They share the same idea of decomposing the adversarial robustness like TRADES and utilize unlabeled data for stability while labeled data for classification. Empirically, they investigated the impact of different factors on adversarial training like label noise, distribution shift, and the amount of additional data. On the other hand, Najafi *et al.* (2019) introduced some new complexity measures like Adversarial Rademacher Complexity and Minimum Supervision Ratio for theoretical analysis on generalization. It is also observed that adversarial robustness is benefited by self-supervised training [Hendrycks *et al.*, 2019].

It is inspiring to see the improvement of adversarial robustness brought by additional unlabeled data. However, theoretical or empirical guarantees that how much additional data are needed precisely still lack. Besides, the cost of such methods should not be neglected, including collecting data and training adversarially on data multiple times larger than original datasets.

Efficient Adversarial Training

One well-known limitation of conventional adversarial training methods like PGD-AT is that their training time is 3-30 times longer than that in standard training [Shafahi *et al.*, 2019]. The main reason is the min-max problem described in Equation (3) is solved iteratively. This line of research of adversarial training aims to reduce the time cost while keeping the performances of adversarial training.

As the first attempt, the core idea of free adversarial training (Free-AT) [Shafahi *et al.*, 2019] is to reuse the gradients computed in the backward pass when doing forward pass. In Free-AT, both the model parameters and image perturbations are updated simultaneously. Concretely, for the same mini-batch data, the same operation is done for m times in a row, equivalent to utilizing strong adversarial examples in PGD-AT. Further upon Free-AT, Wong *et al.* (2020) proposed fast adversarial training (FAST-AT), which utilizes FGSM with random initialization and is as effective as the PGD-AT. They also attributed the failure of FGSM-based adversarial training methods to the *catastrophic overfitting* and zero-initialized perturbation. However, Andriushchenko and Flammarion (2020) found these fast training methods [Shafahi *et al.*, 2019; Wong *et al.*, 2020] suffer from catastrophic overfitting as well. They also pointed out the reason for randomization to take effect in [Wong *et al.*, 2020] is that randomization slightly reduces the magnitude of perturbations. Kim *et al.* (2021) supported the above finding and demonstrated catastrophic overfitting is because single-step adversarial training uses only adversarial examples with maximum perturbations. For the purpose of preventing catastrophic overfitting, many improvements are proposed like GradAlign [Andriushchenko and Flammarion, 2020], dynamic schedule [S. and Babu, 2020], inner inter-

val verification [Kim *et al.*, 2021], domain adaption [Song *et al.*, 2019] and regularization methods [Vivek and Babu, 2020; Huang *et al.*, 2020].

Intrinsic from the above works, Zhang *et al.* (2019a) proposed You Only Propagate Once (YOPO) from the perspective of Pontryagin’s Maximum Principle. According to their analysis on adversarial training, they discovered that adversarial gradients update is only related to the first layer of neural networks. This property enables YOPO to focus on the first layer of the proposed network architecture for adversary computation while other layers are frozen, significantly reducing the numbers of forward and backward propagation. The authors claimed that Free-AT is a particular case of YOPO.

Other Variants

In addition to the above branches of adversarial training methods, several other variants of adversarial training are summarized as follows. Some works modify the learning objectives of vanilla adversarial training, like adversarial distributional training [Dong *et al.*, 2020] where a distribution-based min-max problem is derived from a general view; bilateral adversarial training [Wang and Zhang, 2019] where the model is training on both perturbed images and labels; adversarial training based on feature scatter [Zhang and Wang, 2019], which utilizes a distance metric for sets of natural data and their counterparts and produces adversarial examples in feature space; and adversarial training integrated with self-supervised learning [Kim *et al.*, 2020; Jiang *et al.*, 2020]. Some replace the fundamental components of models for better performances, like hypersphere embedding [Pang *et al.*, 2020], and smoothed ReLU function [Xie *et al.*, 2020]. Lastly, Lee *et al.* (2020) proposed AVmixup to augment adversarial examples by interpolation.

3.3 Generalization Problem in Adversarial Training

For deep learning algorithms, generalization is a significant characteristic. Though most efforts in the research community are paid for improving adversarial training under given adversarial attacks, the voice of discussions on generalization is getting louder. This subsection mainly reviews the studies on the generalization of adversarial training from three aspects: standard generalization, adversarially robust generalization, and generalization on unseen attacks.

Standard Generalization

Despite the success in improving the robustness of neural networks to adversarial attacks, adversarial training is observed that hurts standard accuracy badly [Madry *et al.*, 2018], leading to the discussion of relationships between adversarial robustness and standard accuracy. We refer to it as standard generalization.

One popular viewpoint is the trade-off between adversarial robustness and standard accuracy. Tsipras *et al.* (2019) claimed that standard accuracy and adversarial robustness might at odds and demonstrated the existence of the trade-off via a binary classification task. Su *et al.* (2018) evaluated the recent SOTA ImageNet-based models on multiple robustness metrics. They concluded a linearly negative correlation

between the logarithm of model classification accuracy and model robustness. Zhang *et al.* (2019b) decomposed the robust error as the sum of natural error and the boundary error and provided a tight upper bound for them, theoretically characterizing the trade-off.

However, some works have different opinions that adversarial robustness and standard accuracy are not opposing. Stutz *et al.* (2019) studied the manifold of adversarial examples and natural data. They confirmed the existence of adversarial examples on the manifold of natural data, adversarial robustness on which is equivalent to generalization. Yang *et al.* (2020b) investigated various datasets, from MNIST to Restricted ImageNet, showing these datasets are distributionally separated, and the separation is usually larger than 2ϵ (the value of ϵ differs in different datasets, *e.g.*, 0.1/1 for MNIST, 8/255 for CIFAR). It indicates the existence of robust and accurate classifiers. They also claimed existing training methods fail to impose local Lipschitzness or are insufficiently generalized. Experiments in [Raghunathan *et al.*, 2019] support this statement, where additional unlabeled data is proved to help mitigate the trade-off.

As suggested in [Yang *et al.*, 2020b], the trade-off might not be inherent but a consequence of current adversarial training methods. Though researchers haven't reached a consensus on the cause of the trade-off, existing evidence does reveal some limitations on adversarial training. Adversarial robustness should not be at the cost of standard accuracy. Some variants of adversarial training show better standard generalization empirically, such as adaptive ϵ for adversarial training reviewed in Section 3.2, robust local features [Song *et al.*, 2020] and \mathcal{L}_1 penalty [Xing *et al.*, 2020].

Adversarially Robust Generalization

The phenomenon that adversarially trained models do not perform well on adversarially perturbed test data is firstly observed in [Madry *et al.*, 2018]. In other words, there is a large gap between the training accuracy and test accuracy on adversarial data. Other than CIFAR-10, similar experimental results are observed on multiple datasets, such as SVHN, CIFAR-100, and ImageNet [Rice *et al.*, 2020]. These gaps indicate that a severe overfitting happens in current adversarial training methods. Such overfitting is initially studied by [Schmidt *et al.*, 2018], who refer to it as adversarially robust generalization.

Schmidt *et al.* (2018) revealed the difficulty of obtaining a robust model with the fact that more training data are required for adversarially robust generalization. Later many efforts are made to improve generalization empirically, such as adversarial training with semi/unsupervised Learning, AVmixup, and robust local feature [Song *et al.*, 2020]. In contrast, Rice *et al.* (2020) systematically investigated various techniques used in deep learning like ℓ_1 and ℓ_2 regularization, cutout, mixup and early stopping, where *early stopping* is found to be the most effective and confirmed by [Pang *et al.*, 2021].

On the other hand, though researchers attempt to analyze this generalization problem with different tools like Rademacher complexity [Yin *et al.*, 2019] and VC dimension [Cullina *et al.*, 2018], theoretical progress is, in fact, limited, and the generalization problem is far from being solved.

Generalization on Unseen Attacks

It is proved that the specific type of attacks are not sufficient to represent the space of possible perturbations [Tramèr *et al.*, 2018; Goodfellow *et al.*, 2015]. However, in adversarial training, the inner maximization problem's constraints: l_p norm and ϵ are pre-fixed. Thus, adversarially trained models, which are robust to a specific attack, *e.g.*, l_∞ adversarial examples, can be circumvented easily by different types of attacks, *e.g.*, other l_p norms, or larger ϵ , or different target models [Kang *et al.*, 2019]. Simply combining perturbations with different l_p norms in adversarial training proves to be useless [Tramèr and Boneh, 2019]. This kind of poor generalization to other attacks significantly degrades the reliability of adversarial training.

Such limitation is intrinsically caused by adversarial training itself, and the key is how to solve the inner problem properly. The research line in EAT can be seen as the first attempt to approximate the optimal solutions to the inner problem by increasing the number and diversity of targeted models during training. Similarly, Maini *et al.* (2020) adopted adversaries under different l_p norm and use the steepest descent to approximate the optimal values for the inner problem. Dong *et al.* (2020) proposed to explicitly model the distribution of adversarial examples around each sample, replacing the the perturbation set $B(x, \epsilon)$. Differently, Stutz *et al.* (2020) attributed the poor generalization to enforcing high-confidence predictions on adversarial examples, and suggested calibrating the confidence scores during adversarial training. Besides, it is worth investigating the connections between adversarial examples and common image corruptions *e.g.*, Gaussian noises [Gilmer *et al.*, 2019]. There are evidences that models trained on such common corruptions show robustness to adversarial examples [Gokhale *et al.*, 2020; Wong and Kolter, 2021]. In the meanwhile, Volpi *et al.* (2018) and Qiao *et al.* (2020) also utilize adversarial data augmentation for better generalization on unseen domains.

Though significant, the generalization problem of adversarial training on unseen attacks is only occasionally studied at this time. One possible reason is that our understanding of adversarial examples is limited and incomplete. The truth of adversarial examples is still underground, which also needs much effort.

4 Conclusion and Future Directions

In this paper, we present recent advances of adversarial training methods for adversarial robustness. To our best knowledge, for the first time, we review adversarial training with a novel taxonomy and discuss the generalization problem in adversarial training. We also summarize the benchmarks and provide performance comparisons of different methods. Despite extensive efforts, the vulnerability of deep learning models to adversarial examples has not been completely solved by adversarial training. Several open problems remain yet to solve, which are summarized as follows.

Min-Max Optimization in Adversarial Training. Adversarial training is formulated as a min-max problem. However, due to the non-convexity of deep neural networks, it is very challenging to obtain the global optimum for adversarial

training. In existing methods, PGD is a prevalent technique for approximating the optimum, as Madry *et al.* (2018) empirically proved the tractability of adversarial training with PGD. But it can hardly provide a “robustness certificate” after solving the problem [Razaviyayn *et al.*, 2020]. In other words, the robustness of adversarially trained models is not guaranteed [Kang *et al.*, 2019]. For this purpose, the development of new techniques for solving non-convex min-max problems is necessary and crucial.

Overfitting in Adversarial Training. Overfitting is a common problem in deep learning, and there are effective countermeasures to alleviate the overfitting. In adversarial training, however, overfitting seems to be more severe and those widely-used techniques in deep learning help little [Rice *et al.*, 2020]. The generalization gap between adversarial training accuracy and testing accuracy is very large. From the perspective of generalization, the theory of sample complexity [Schmidt *et al.*, 2018] explains such a phenomenon partially, and is supported by experimental results in derivative works [Alayrac *et al.*, 2019; Carmon *et al.*, 2019]. As suggested by [Schmidt *et al.*, 2018], it is essential to explore the intersections between robustness, classifiers and data distribution. Some open problems can be found in [Schmidt *et al.*, 2018].

Beyond Adversarial Training. Though many theories have been proposed for improving adversarial training, it is undeniable that these improvements are less effective than claimed [Pang *et al.*, 2021]. Some basic settings *e.g.*, training schedule, early stopping, seem to owe much on the improvement of adversarial training. The shreds of evidence in [Yang *et al.*, 2020b; Stutz *et al.*, 2019] show that adversarial training might not be the optimal solution for obtaining models with desirable robustness and accuracy. The trade-off between robustness and generalization can also be seen as an intrinsic limitation of adversarial training. Thus it is critical and necessary to investigate new methods beyond adversarial training for adversarial robustness in the future.

Adversarial Training in Other Domains. This paper discusses the recent advances of adversarial training in the image domain, which is the most active research field of adversarial robustness. However, adversarial training is not just limited to images. Adversarial training has been successfully applied to texts [Miyato *et al.*, 2017], graphs [Dai *et al.*, 2019], audios [Pandey and Wang, 2018], and reinforcement learning [Pattanaik *et al.*, 2017]. Therefore, we hope this survey can provide helpful insights to interested researchers and inspire more progress in different domains.

Acknowledgements

This paper is supported by 1) Singapore Ministry of Education Academic Research Fund Tier 1 RG128/18, Tier 1 RG115/19, Tier 1 RT07/19, Tier 1 RT01/19, Tier 1 RG24/20, and Tier 2 MOE2019-T2-1-176, 2) NTU-WASP Joint Project, 3) Singapore NRF National Satellite of Excellence, Design Science and Technology for Secure Critical Infrastructure NSoE DeST-SCI2019-0012, 4) AI Singapore (AISG) 100 Experiments (100E) programme, and 5) NTU

Project for Large Vertical Take-Off & Landing (VTOL) Research Platform. Bihan Wen was supported in part by the National Research Foundation (NRF), Singapore, through the Singapore Cybersecurity Consortium (SGCSC) Grant Office, under SGCSC_Grant_2019-S01. Qian Wang’s work was partially supported by the NSFC under Grants U20B2049 and 61822207. We would like to express our very great appreciation to Prof Bo Li and Mr. Qi Bi for their valuable and constructive suggestions of this research work.

References

- [Alayrac *et al.*, 2019] Jean-Baptiste Alayrac, Jonathan Uesato, Po-Sen Huang, Alhussein Fawzi, Robert Stanforth, and Pushmeet Kohli. Are labels required for improving adversarial robustness? In *NeurIPS*, pages 12192–12202, 2019.
- [Andriushchenko and Flammarion, 2020] Maksym Andriushchenko and Nicolas Flammarion. Understanding and improving fast adversarial training. In *NeurIPS*, 2020.
- [Athalye and Carlini, 2018] Anish Athalye and Nicholas Carlini. On the robustness of the cvpr 2018 white-box adversarial example defenses. *arXiv preprint arXiv:1804.03286*, 2018.
- [Balaji *et al.*, 2019] Yogesh Balaji, Tom Goldstein, and Judy Hoffman. Instance adaptive adversarial training: Improved accuracy tradeoffs in neural nets. *arXiv preprint arXiv:1910.08051*, 2019.
- [Buckman *et al.*, 2018] Jacob Buckman, Aurko Roy, Colin Raffel, and Ian J. Goodfellow. Thermometer encoding: One hot way to resist adversarial examples. In *Proc. of ICLR*, 2018.
- [Cai *et al.*, 2018] Qi-Zhi Cai, Chang Liu, and Dawn Song. Curriculum adversarial training. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 3740–3747, 2018.
- [Carmon *et al.*, 2019] Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C. Duchi, and Percy Liang. Unlabeled data improves adversarial robustness. In *NeurIPS*, pages 11190–11201, 2019.
- [Cheng *et al.*, 2020] Minhao Cheng, Qi Lei, Pin-Yu Chen, Inderjit Dhillon, and Cho-Jui Hsieh. Cat: Customized adversarial training for improved robustness. *arXiv preprint arXiv:2002.06789*, 2020.
- [Cullina *et al.*, 2018] Daniel Cullina, Arjun Nitin Bhagoji, and Prateek Mittal. Pac-learning in the presence of evasion adversaries. *arXiv preprint arXiv:1806.01471*, 2018.
- [Dai *et al.*, 2019] Quanyu Dai, Xiao Shen, Liang Zhang, Qiang Li, and Dan Wang. Adversarial training methods for network embedding. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pages 329–339, 2019.
- [Dalvi *et al.*, 2004] Nilesh Dalvi, Pedro Domingos, Sumitanghai, and Deepak Verma. Adversarial classification. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2004.

- [Dauphin *et al.*, 2014] Yann N. Dauphin, Razvan Pascanu, Çağlar Gülçehre, KyungHyun Cho, Surya Ganguli, and Yoshua Bengio. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *NIPS*, pages 2933–2941, 2014.
- [Ding *et al.*, 2020] Gavin Weiguang Ding, Yash Sharma, Kry Yik Chau Lui, and Ruitong Huang. MMA training: Direct input space margin maximization through adversarial training. In *Proc. of ICLR*, 2020.
- [Dong *et al.*, 2020] Yinpeng Dong, Zhijie Deng, Tianyu Pang, Jun Zhu, and Hang Su. Adversarial distributional training for robust deep learning. In *NeurIPS*, 2020.
- [Engstrom *et al.*, 2018] Logan Engstrom, Andrew Ilyas, and Anish Athalye. Evaluating and understanding the robustness of adversarial logit pairing. *arXiv preprint arXiv:1807.10272*, 2018.
- [Gilmer *et al.*, 2019] Justin Gilmer, Nicolas Ford, Nicholas Carlini, and Ekin D. Cubuk. Adversarial examples are a natural consequence of test error in noise. In *ICML*, pages 2280–2289, 2019.
- [Gokhale *et al.*, 2020] Tejas Gokhale, Rushil Anirudh, Bhavya Kailkhura, Jayaraman J Thiagarajan, Chitta Baral, and Yezhou Yang. Attribute-guided adversarial training for robustness to natural perturbations. *arXiv preprint arXiv:2012.01806*, 2020.
- [Goodfellow *et al.*, 2015] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *Proc. of ICLR*, 2015.
- [Hendrycks *et al.*, 2019] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. In *NeurIPS*, pages 15637–15648, 2019.
- [Huang *et al.*, 2015] Ruitong Huang, Bing Xu, Dale Schuurmans, and Csaba Szepesvári. Learning with a strong adversary. *arXiv preprint arXiv:1511.03034*, 2015.
- [Huang *et al.*, 2020] Tianjin Huang, Vlado Menkovski, Yulong Pei, and Mykola Pechenizkiy. Bridging the performance gap between fgm and pgd adversarial training. *arXiv preprint arXiv:2011.05157*, 2020.
- [Ilyas *et al.*, 2019] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. In *NeurIPS*, pages 125–136, 2019.
- [Jiang *et al.*, 2020] Ziyu Jiang, Tianlong Chen, Ting Chen, and Zhangyang Wang. Robust pre-training by adversarial contrastive learning. In *NeurIPS*, 2020.
- [Kang *et al.*, 2019] Daniel Kang, Yi Sun, Tom Brown, Dan Hendrycks, and Jacob Steinhardt. Transfer of adversarial robustness between perturbation types. *arXiv preprint arXiv:1905.01034*, 2019.
- [Kannan *et al.*, 2018] Harini Kannan, Alexey Kurakin, and Ian Goodfellow. Adversarial logit pairing. *arXiv preprint arXiv:1803.06373*, 2018.
- [Kariyappa and Qureshi, 2019] Sanjay Kariyappa and Moinuddin K Qureshi. Improving adversarial robustness of ensembles with diversity training. *arXiv preprint arXiv:1901.09981*, 2019.
- [Kim *et al.*, 2020] Minseon Kim, Jihoon Tack, and Sung Ju Hwang. Adversarial self-supervised contrastive learning. In *NeurIPS*, 2020.
- [Kim *et al.*, 2021] Hoki Kim, Woojin Lee, and Jaewook Lee. Understanding catastrophic overfitting in single-step adversarial training. In *Proceedings of AAAI 2021*, 2021.
- [Kurakin *et al.*, 2016] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*, 2016.
- [Kurakin *et al.*, 2017] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial machine learning at scale. In *Proc. of ICLR*, 2017.
- [Lee *et al.*, 2020] Saehyung Lee, Hyungyu Lee, and Sungroh Yoon. Adversarial vertex mixup: Toward better adversarially robust generalization. In *IEEE/CVF CVPR*, pages 269–278, 2020.
- [Lowd and Meek, 2005] Daniel Lowd and Christopher Meek. Adversarial learning. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, 2005.
- [Madry *et al.*, 2018] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *Proc. of ICLR*, 2018.
- [Maini *et al.*, 2020] Pratyush Maini, Eric Wong, and J. Zico Kolter. Adversarial robustness against the union of multiple perturbation models. In *ICML*, pages 6640–6650, 2020.
- [Mao *et al.*, 2019] Chengzhi Mao, Ziyuan Zhong, Junfeng Yang, Carl Vondrick, and Baishakhi Ray. Metric learning for adversarial robustness. In *NeurIPS*, pages 478–489, 2019.
- [Miyato *et al.*, 2017] Takeru Miyato, Andrew M. Dai, and Ian J. Goodfellow. Adversarial training methods for semi-supervised text classification. In *Proc. of ICLR*, 2017.
- [Najafi *et al.*, 2019] Amir Najafi, Shin-ichi Maeda, Masanori Koyama, and Takeru Miyato. Robustness to adversarial perturbations in learning from incomplete data. In *NeurIPS*, pages 5542–5552, 2019.
- [Pandey and Wang, 2018] Ashutosh Pandey and DeLiang Wang. On adversarial training and loss functions for speech enhancement. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2018, Calgary, AB, Canada, April 15-20, 2018*, pages 5414–5418, 2018.
- [Pang *et al.*, 2019] Tianyu Pang, Kun Xu, Chao Du, Ning Chen, and Jun Zhu. Improving adversarial robustness via promoting ensemble diversity. In *ICML*, pages 4970–4979, 2019.

- [Pang *et al.*, 2020] Tianyu Pang, Xiao Yang, Yinpeng Dong, Kun Xu, Jun Zhu, and Hang Su. Boosting Adversarial Training with Hypersphere Embedding. In *NeurIPS*, pages 7779–7792, 2020.
- [Pang *et al.*, 2021] Tianyu Pang, Xiao Yang, Yinpeng Dong, Hang Su, and Jun Zhu. Bag of tricks for adversarial training. In *International Conference on Learning Representations*, 2021.
- [Papernot *et al.*, 2017] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z. Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. *ASIA CCS 2017*, 2017.
- [Pattanaik *et al.*, 2017] Anay Pattanaik, Zhenyi Tang, Shuijing Liu, Gautham Bommanna, and Girish Chowdhary. Robust deep reinforcement learning with adversarial attacks. *arXiv preprint arXiv:1712.03632*, 2017.
- [Qiao *et al.*, 2020] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *IEEE/CVF CVPR*, pages 12553–12562, 2020.
- [Qin *et al.*, 2019] Chongli Qin, James Martens, Sven Gowal, Dilip Krishnan, Krishnamurthy Dvijotham, Alhussein Fawzi, Soham De, Robert Stanforth, and Pushmeet Kohli. Adversarial robustness through local linearization. In *NeurIPS*, pages 13824–13833, 2019.
- [Raghunathan *et al.*, 2019] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John C Duchi, and Percy Liang. Adversarial training can hurt generalization. *arXiv preprint arXiv:1906.06032*, 2019.
- [Razaviyayn *et al.*, 2020] Meisam Razaviyayn, Tianjian Huang, Songtao Lu, Maher Nouiehed, Maziar Sanjabi, and Mingyi Hong. Nonconvex min-max optimization: Applications, challenges, and recent theoretical advances. *IEEE Signal Processing Magazine*, (5), 2020.
- [Rice *et al.*, 2020] Leslie Rice, Eric Wong, and J. Zico Kolter. Overfitting in adversarially robust deep learning. In *ICML*, pages 8093–8104, 2020.
- [S. and Babu, 2020] Vivek B. S. and R. Venkatesh Babu. Single-step adversarial training with dropout scheduling. In *IEEE/CVF CVPR*, pages 947–956, 2020.
- [Sadeghi *et al.*, 2020] Koosha Sadeghi, Ayan Banerjee, and Sandeep K. S. Gupta. A system-driven taxonomy of attacks and defenses in adversarial machine learning. *IEEE Transactions on Emerging Topics in Computational Intelligence*, (4), 2020.
- [Schmidt *et al.*, 2018] Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, and Aleksander Madry. Adversarially robust generalization requires more data. In *NeurIPS*, pages 5019–5031, 2018.
- [Schott *et al.*, 2019] Lukas Schott, Jonas Rauber, Matthias Bethge, and Wieland Brendel. Towards the first adversarially robust neural network model on MNIST. In *Proc. of ICLR*, 2019.
- [Shafahi *et al.*, 2019] Ali Shafahi, Mahyar Najibi, Amin Ghiasi, Zheng Xu, John P. Dickerson, Christoph Studer, Larry S. Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! In *NeurIPS*, pages 3353–3364, 2019.
- [Shaham *et al.*, 2018] Uri Shaham, Yutaro Yamada, and Sahand Negahban. Understanding adversarial training: Increasing local stability of supervised models through robust optimization. *Neurocomputing*, 2018.
- [Silva and Najafirad, 2020] Samuel Henrique Silva and Peyman Najafirad. Opportunities and Challenges in Deep Learning Adversarial Robustness: A Survey. *arXiv preprint arXiv:2007.00753*, 2020.
- [Smith, 2019] Leslie N Smith. A useful taxonomy for adversarial robustness of neural networks. *arXiv preprint arXiv:1910.10679*, 2019.
- [Song *et al.*, 2019] Chuanbiao Song, Kun He, Liwei Wang, and John E. Hopcroft. Improving the generalization of adversarial training with domain adaptation. In *Proc. of ICLR*, 2019.
- [Song *et al.*, 2020] Chuanbiao Song, Kun He, Jiadong Lin, Liwei Wang, and John E. Hopcroft. Robust local features for improving the generalization of adversarial training. In *Proc. of ICLR*, 2020.
- [Stutz *et al.*, 2019] David Stutz, Matthias Hein, and Bernt Schiele. Disentangling adversarial robustness and generalization. In *CVPR*, pages 6976–6987, 2019.
- [Stutz *et al.*, 2020] David Stutz, Matthias Hein, and Bernt Schiele. Confidence-calibrated adversarial training: Generalizing to unseen attacks. In *ICML*, pages 9155–9166, 2020.
- [Su *et al.*, 2018] Dong Su, Huan Zhang, Hongge Chen, Jinfeng Yi, Pin-Yu Chen, and Yupeng Gao. Is robustness the cost of accuracy?—a comprehensive study on the robustness of 18 deep image classification models. In *ECCV*, pages 631–648, 2018.
- [Szegedy *et al.*, 2014] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *Proc. of ICLR*, 2014.
- [Tabassi *et al.*, 2019] Elham Tabassi, Kevin Burns, Michael Hadjimichael, Andres Molina-Markham, and Julian Sexton. A taxonomy and terminology of adversarial machine learning. *NIST IR*, 2019.
- [Tramèr and Boneh, 2019] Florian Tramèr and Dan Boneh. Adversarial training and robustness for multiple perturbations. In *NeurIPS*, pages 5858–5868, 2019.
- [Tramèr *et al.*, 2017] Florian Tramèr, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. The space of transferable adversarial examples. *arXiv preprint arXiv:1704.03453*, 2017.
- [Tramèr *et al.*, 2018] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian J. Goodfellow, Dan Boneh, and Patrick D. McDaniel. Ensemble adversarial training: Attacks and defenses. In *Proc. of ICLR*, 2018.

- [Tsipras *et al.*, 2019] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. In *Proc. of ICLR*, 2019.
- [Vivek and Babu, 2020] BS Vivek and R Venkatesh Babu. Regularizers for single-step adversarial training. *arXiv preprint arXiv:2002.00614*, 2020.
- [Volpi *et al.*, 2018] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C. Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. In *NeurIPS*, pages 5339–5349, 2018.
- [Wang and Zhang, 2019] Jianyu Wang and Haichao Zhang. Bilateral adversarial training: Towards fast training of more robust models against adversarial attacks. In *ICCV*, pages 6628–6637, 2019.
- [Wang *et al.*, 2019] Yisen Wang, Xingjun Ma, James Bailey, Jinfeng Yi, Bowen Zhou, and Quanquan Gu. On the convergence and robustness of adversarial training. In *ICML*, pages 6586–6595, 2019.
- [Wang *et al.*, 2020] Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. Improving adversarial robustness requires revisiting misclassified examples. In *Proc. of ICLR*, 2020.
- [Wong and Kolter, 2021] Eric Wong and J Zico Kolter. Learning perturbation sets for robust machine learning. In *Proc. of ICLR*, 2021.
- [Wong *et al.*, 2020] Eric Wong, Leslie Rice, and J. Zico Kolter. Fast is better than free: Revisiting adversarial training. In *Proc. of ICLR*, 2020.
- [Xie *et al.*, 2020] Cihang Xie, Mingxing Tan, Boqing Gong, Alan Yuille, and Quoc V Le. Smooth adversarial training. *arXiv preprint arXiv:2006.14536*, 2020.
- [Xing *et al.*, 2020] Yue Xing, Qifan Song, and Guang Cheng. On the Generalization Properties of Adversarial Training. *arXiv preprint arXiv:2008.06631*, 2020.
- [Yang *et al.*, 2020a] Huanrui Yang, Jingyang Zhang, Hongliang Dong, Nathan Inkawhich, Andrew Gardner, Andrew Touchet, Wesley Wilkes, Heath Berry, and Hai Li. DVERGE: Diversifying Vulnerabilities for Enhanced Robust Generation of Ensembles. In *NeurIPS*, 2020.
- [Yang *et al.*, 2020b] Yao-Yuan Yang, Cyrus Rashtchian, Hongyang Zhang, Ruslan Salakhutdinov, and Kamalika Chaudhuri. A closer look at accuracy vs. robustness. *Advances in Neural Information Processing Systems*, 33, 2020.
- [Yin *et al.*, 2019] Dong Yin, Kannan Ramchandran, and Peter L. Bartlett. Rademacher complexity for adversarially robust generalization. In *ICML*, pages 7085–7094, 2019.
- [Yuan *et al.*, 2019] Xiaoyong Yuan, Pan He, Qile Zhu, and Xiaolin Li. Adversarial Examples: Attacks and Defenses for Deep Learning. *IEEE TNNLS*, 2019.
- [Zhai *et al.*, 2019] Runtian Zhai, Tianle Cai, Di He, Chen Dan, Kun He, John Hopcroft, and Liwei Wang. Adversarially robust generalization just requires more unlabeled data. *arXiv preprint arXiv:1906.00555*, 2019.
- [Zhang and Wang, 2019] Haichao Zhang and Jianyu Wang. Defense against adversarial attacks using feature scattering-based adversarial training. In *NeurIPS*, pages 1829–1839, 2019.
- [Zhang *et al.*, 2019a] Dinghuai Zhang, Tianyuan Zhang, Yiping Lu, Zhanxing Zhu, and Bin Dong. You only propagate once: Accelerating adversarial training via maximal principle. In *NeurIPS*, pages 227–238, 2019.
- [Zhang *et al.*, 2019b] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P. Xing, Laurent El Ghaoui, and Michael I. Jordan. Theoretically principled trade-off between robustness and accuracy. In *ICML*, pages 7472–7482, 2019.
- [Zhang *et al.*, 2020] Jinfeng 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 *ICML*, 2020.