



Pareto Adversarial Robustness: Balancing Spatial Robustness and Sensitivity-based Robustness

Journal:	SCIENCE CHINA Information Sciences
Manuscript ID	SCIS-2022-1210.R3
Manuscript Type:	Research Paper
Date Submitted by the Author:	01-Sep-2023
Complete List of Authors:	Sun, Ke; Peking University, Li, Mingjie Lin, Zhouchen; Peking University, Department of Intelligence Sciences
Keywords:	Deep Learning, Adversarial Robustness, Reliable Machine Learning, Pareto Optimization, Spatial Robustness
Specialty:	Artificial Intelligence\Pattern Recognition\Database and Knowledge-Based Systems < Computer Science & Technology

SCHOLARONE™
Manuscripts

• RESEARCH PAPER •

Pareto Adversarial Robustness: Balancing Spatial
Robustness and Sensitivity-based Robustness

Ke Sun¹, Mingjie Li¹ & Zhouchen Lin^{1,2,3*}

¹Key Lab. of Machine Perception (MoE), School of Intelligence Science and Technology, Peking University, Beijing 100871, China;

²Institute for Artificial Intelligence, Peking University, Beijing 100871, China;

³Peng Cheng Laboratory, Guangzhou 510335, China

Abstract Adversarial robustness, which primarily comprises sensitivity-based robustness and spatial robustness, plays an integral part in achieving robust generalization. In this paper, we endeavor to design strategies to achieve universal adversarial robustness. To achieve this, we first investigate the relatively less-explored realm of spatial robustness. Then, we integrate the existing spatial robustness methods by incorporating both local and global spatial vulnerability into a unified spatial attack and adversarial training approach. Furthermore, we present a comprehensive relationship between natural accuracy, sensitivity-based robustness, and spatial robustness, supported by strong evidence from the perspective of robust representation. Crucially, to reconcile the interplay between the mutual impacts of various robustness components into one unified framework, we incorporate the *Pareto criterion* into the adversarial robustness analysis, yielding a novel strategy called *Pareto Adversarial Training* for achieving universal robustness. The resulting Pareto front, which delineates the set of optimal solutions, provides an optimal balance between natural accuracy and various adversarial robustness. This sheds light on solutions for achieving universal robustness in the future. To the best of our knowledge, we are the first to consider universal adversarial robustness via multi-objective optimization.

Keywords Deep Learning, Adversarial Robustness, Reliable Machine Learning, Pareto Optimization, Spatial Robustness

Citation

1 Introduction

Robust generalization serves as an extension of the traditional generalization that is normally achieved via Empirical Risk Minimization for i.i.d. data [33]. However, the test environment could be slightly or dramatically different from the training environment [15] in a robust generalization scenario. Lately, improving the robustness of deep neural networks has been one of the pivotal areas of research, encompassing different threads of research such as adversarial robustness [8, 28], non-adversarial robustness [10, 37], Bayesian deep learning [7, 23] and causality [1]. In this paper, we focus on adversarial robustness, where adversarial examples are carefully manipulated by humans to fool machine learning models, e.g., deep neural networks, which could pose serious threats, especially in safety-critical applications. Currently, adversarial training [4, 8, 21, 36] is regarded as a promising and widely accepted strategy to address this issue.

Like Out-of-Distribution (OoD) robustness, adversarial robustness also has several aspects [3, 9, 20], including *sensitivity-based robustness* [30], i.e., robustness against pixel-wise perturbations (normally within the constraints of an l_p ball), and *spatial robustness*, i.e., robustness against multiple spatial transformations. Computer vision and graphics literature provide a deeper insight into these two aspects, revealing that two main factors determine the appearance of a pictured object [29, 35]: (1) lighting and materials, and (2) geometry. Most previous studies on adversarial robustness have focused only on the first factor [35] by examining pixel-wise perturbations, e.g., Projected Gradient Descent (PGD) attacks [21],

* Corresponding author (email: zlin@pku.edu.cn)

assuming that the underlying geometry stays the same after the adversarial perturbation. Only a small proportion of research works have attempted to tackle the less-studied second factor, which includes Flow-based [35] and Rotation-Translation (RT)-based attacks [5, 6].

However, it is crucial to consider spatial robustness for achieving universal robustness, the ultimate objective of robust generalization. One of the most important reasons is that sensitivity-based robustness, which is generally based on the l_p -distance, is not sufficient to maintain perceptual similarity [5, 6, 25, 35]. Specifically, although *spatial attacks or geometric transformations* result in small perceptual differences, they yield large l_p distances.

A clear relationship between accuracy, sensitivity-based and spatial robustness is the key to achieving universal adversarial robustness. While the trade-off between *sensitivity-based robustness* and accuracy has been revealed by several studies [24, 32, 40], the comprehensive relationships among *spatial robustness* and them are still unclear. Although previous studies [12, 31] have explored this issue, they only focused on Rotation-Translation spatial robustness and did not consider Flow-based spatial robustness [35, 39]. Surprisingly, we find that Flow-based spatial robustness presents a relationship contrary to the one revealed previously, making the previous conclusion less reliable.

Based on this important finding, we start our exploration of clearer relationships between different robustnesses, and we eventually harmonize the conflicting relationships within them by leveraging the Pareto criterion [13, 14, 38], thus achieving an optimal balanced universal robustness. A recent study [24] attributes the conflicting relationships among the various robustnesses to overparametrization, while we uncovered it from the perspective of different shape-biased representations. Another report [34] examined the trade-off in the inference time, while we target more comprehensive relationships between different robustnesses with a different methodology.

In this paper, we first try to gain deeper insights into the robustness relationships by investigating the two main spatial robustness branches, i.e., Flow-based spatial attack [35] and Rotation-Translation (RT) attack [5]. After revealing their impact on local and global spatial sensitivity, we propose integrated spatial attack and spatial adversarial training, which can incorporate comprehensive spatial vulnerabilities or robustness. Based on this understanding, we present a comprehensive relationship among the accuracy, sensitivity-based robustness, and the two branches of spatial robustness by investigating their different saliency maps from the perspectives of shape-bias, sparse or dense representation. It turns out that while the relationship between sensitivity-based and RT robustness is a fundamental trade-off, sensitivity-based and Flow-based spatial robustness are highly correlated, providing a vital supplementary for previous conclusions. Thus, comprehensive relationships between accuracy and the various robustnesses are not pure trade-offs, motivating us to introduce the Pareto criterion [13, 14, 38], the general multi-objective optimization principle, into the universal adversarial robustness analysis. The Pareto criterion enables an optimal balance between the interplay of natural accuracy and the different adversarial robustnesses, leading to universal adversarial robustness in a Pareto manner. By incorporating a two-moment term that can capture the interaction between loss of accuracy and different robustnesses, we propose a bi-level optimization framework called *Pareto Adversarial Training*. The resulting Pareto front provides a set of optimal solutions that can balance perfectly all the relationships under consideration, outperforming other existing strategies.

Our contributions can be summarized as follows:

- We reveal the existence of both local and global spatial robustness and propose integrated spatial attack and spatial adversarial training, incorporating comprehensive spatial vulnerabilities.
- We present comprehensive relationships among accuracy, sensitivity-based, and different spatial robustnesses, supported by strong and intuitive evidence from the perspective of robust representation.
- We incorporate the Pareto criterion into adversarial robustness analysis, and the resulting Pareto Adversarial Training can optimally balance multiple adversarial robustness, yielding universal adversarial robustness.

2 Local and Global Spatial Robustness

To present the comprehensive relationships between accuracy and different adversarial robustnesses, we first provide a fine-grained understanding of spatial robustness. We summarize several studies about spatial robustness [5, 6, 12, 31, 35, 39] into two major branches: (1) Flow-based Attacks, and (2) Rotation-Translation (RT) Attacks. In particular, we find that the former mainly focuses on the local spatial

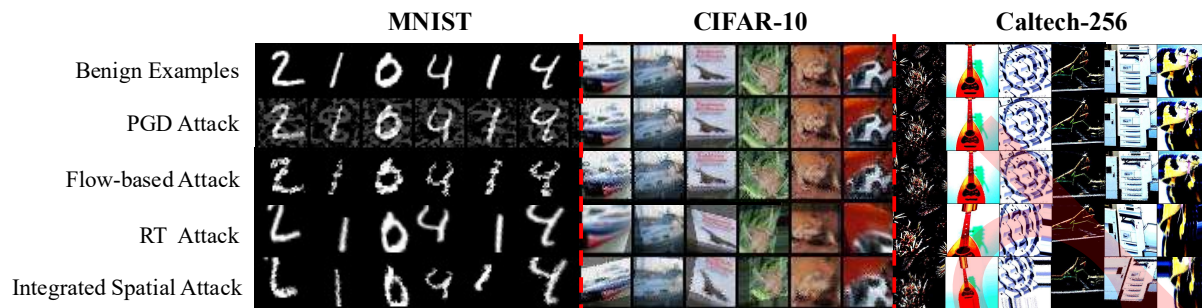


Figure 1 Visualization of Flow-based, RT and Our Integrated Spatial adversarial examples on MNIST, CIFAR-10 and Caltech-256. More images and detailed discussions are provided in Appendix A.

vulnerability while the latter tends to capture the global spatial sensitivity. Based on this finding, integrated spatial attack and spatial adversarial training are proposed.

2.1 Local Spatial Robustness: Flow-based Attacks

The most representative Flow-based Attack is the Spatial Transformed Attack [35], wherein a differentiable flow vector $w_F = (\Delta\mu, \Delta v)$ is introduced in the 2D coordinates (μ, v) to craft adversarial spatial transformation. The vanilla-targeted Flow-based attack [35] follows the optimization manner ($\kappa = 0$):

$$w_F^* = \arg \min_{w_F} \max_{i \neq t} f_{\theta}^i(x_{w_F}) - f_{\theta}^t(x_{w_F}) + \tau \mathcal{L}_{flow}(w_F), \quad (1)$$

where $f_{\theta}(x) = (f_{\theta}^1(x), \dots, f_{\theta}^K(x))$ is the classifier in the K -classification task. x_{w_F} is a Flow-based adversarial example parameterized by the flow vector w_F . \mathcal{L}_{flow} , which measures the local smoothness of the spatial transformation balanced by τ .

Interestingly, our empirical study shown in the left part of Figure 1 suggests that the Flow-based attack tends to yield local permutations among pixels in some specific regions, irrespective of the option of τ , rather than a global spatial transformation based on their *shapes*. Our analysis indicates that this phenomenon is due to two factors: 1) Local permutations, especially in regions where colors of pixels change dramatically, are already sufficiently sensitive to manipulations, as demonstrated by our empirical results shown in Figure 1. 2) The manner of optimization does not incorporate any sort of *shape transformation* information, e.g., a parametric equation of rotation, as opposed to the vanilla Rotation-Translation attack, which we present in the following. Therefore, we conclude that Flow-based attacks tend to capture the local spatial vulnerability. Further, to design the integrated spatial attack, we transform Eq 1 into its untargeted version under cross-entropy loss with flow vector bounded by an ϵ_F -ball:

$$w_F^* = \arg \max_{w_F} \mathcal{L}_{\theta}^{\text{CE}}(x_{w_F}, y) \quad \text{s.t.} \quad \|w_F\| \leq \epsilon_F, \quad (2)$$

where $\mathcal{L}_{\theta}^{\text{CE}}(x, y) = \log \sum_j \exp(f_{\theta}^j(x)) - f_{\theta}^y(x)$. To maintain a uniform optimization form in our integrated spatial attack, we replace local smoothness term \mathcal{L}_{flow} in Eq. 1 with our familiar l_p constraint and leverage the cross-entropy loss instead of the *max* operation as suggested in [2]. Proposition 1 reveals the correlation between the two losses, indicating that the smooth approximation version of *max* operation in Eq. 1, denoted as \mathcal{L}_{θ}^S , has a parallel updating direction with cross-entropy loss related to w_F . Proof can be found in Appendix B.

Proposition 1. Consider $\mathcal{L}_{\theta}^S(x, y) = \log \sum_{i \neq y} \exp(f_{\theta}^i(x)) - f_{\theta}^y(x)$ as the smooth version loss of Eq. 1 without a local smoothness term. For a fixed (x_{w_F}, y) and θ , we have

$$\nabla_{w_F} \mathcal{L}_{\theta}^{\text{CE}}(x_{w_F}, y) = r(x_{w_F}, y) \nabla_{w_F} \mathcal{L}_{\theta}^S(x_{w_F}, y), \quad (3)$$

where $r(x_{w_F}, y) = \sum_{i \neq y} \exp(f_{\theta}^i(x_{w_F})) / \sum_i \exp(f_{\theta}^i(x_{w_F}))$.

2.2 Global Spatial Robustness: Rotation-Translation Attacks

The original Rotation-Translation attack [5,6] applies parametric equation constraints on 2D coordinates, thus capturing the global spatial information:

$$\begin{bmatrix} u' \\ v' \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} + \begin{bmatrix} \delta u \\ \delta v \end{bmatrix}. \quad (4)$$

To design a generic spatial transformation matrix that can simultaneously consider rotation, translation, cropping, and scaling, we re-parameterize the transform matrix as a generic 6-dimensional affine transformation one, inspired by Spatial Transformer Networks [11]:

$$\begin{bmatrix} u' \\ v' \end{bmatrix} = \left(\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} + \begin{bmatrix} w_{RT}^{11} & w_{RT}^{12} & w_{RT}^{13} \\ w_{RT}^{21} & w_{RT}^{22} & w_{RT}^{23} \end{bmatrix} \right) \begin{bmatrix} u \\ v \\ 1 \end{bmatrix}, \quad (5)$$

where we denote $A_{w_{RT}}$ as the generic 6-dimensional affine transformation matrix, in which each entry of w_{RT} indicates the increment in different spatial aspects. For example, $(w_{RT}^{13}, w_{RT}^{23})$ determines translation. Finally, the optimization form of the resulting generic and differentiable RT-based attack bounded by ϵ_{RT} -ball is expressed as:

$$w_{RT}^* = \arg \max_{w_{RT}} \mathcal{L}_{\theta}^{\text{CE}}(x_{w_{RT}}, y) \quad s.t. \quad \|w_{RT}\| \leq \epsilon_{RT}. \quad (6)$$

2.3 Integrated Spatial Attack

The key to achieving integrated spatial robustness is to design an integrated parameterized sampling grid $\mathcal{T}_{w_{RT}, w_F}(G)$ that can wrap the regular grid with both flow and affine transformation, where G is the generated grid. We show our integrated approach as shown below:

$$\begin{aligned} \mathcal{T}_{w_{RT}, w_F}(G) &= A_{w_{RT}} \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} + w_F, \\ x^{adv} &= \mathcal{T}_{w_{RT}, w_F}(G) \circ x. \end{aligned} \quad (7)$$

Then we sample new x^{adv} by $\mathcal{T}_{w_{RT}, w_F}(G)$ via the differentiable bilinear interpolation [11]. Note that w_F has the same dimensions as the grid G , which are different from the impact of two-dimensional translation parameters in w_{RT} . Then the final loss function of the integrated spatial attack can be presented as:

$$w^* = \arg \max_w \mathcal{L}_{\theta}^{\text{CE}}(x + \eta_w, y), \quad s.t. \quad \|w\| \leq \epsilon, \quad (8)$$

where η_w is the crafted integrated spatial perturbation parameterized by $w = [w_F, w_{RT}]^T$, simultaneously considering both Flow-based and RT spatial sensitivity. Note that η_w itself does not necessarily satisfy the l_p constraint directly. For the implementation, we follow the PGD procedure [21], a common practice in sensitivity-based attacks. We consider the infinity norm of w and different learning rates for the two types of spatial robustness. Therefore, the updating rule of w in each iteration is:

$$\begin{aligned} \begin{bmatrix} \bar{w}_F^{t+1} \\ \bar{w}_{RT}^{t+1} \end{bmatrix} &= \begin{bmatrix} w_F^t \\ w_{RT}^t \end{bmatrix} + \begin{bmatrix} \alpha_F \\ \alpha_{RT} \end{bmatrix} \text{sign}(\nabla_w \mathcal{L}_{\theta}^{\text{CE}}(x_{w^t}, y)), \\ \begin{bmatrix} w_F^{t+1} \\ w_{RT}^{t+1} \end{bmatrix} &= \text{clip}_{\epsilon} \left(\begin{bmatrix} \bar{w}_F^{t+1} \\ \bar{w}_{RT}^{t+1} \end{bmatrix} \right), \\ x_{w^{t+1}}^{t+1} &= \mathcal{T}_{w^{t+1}}(G) \circ x, \end{aligned} \quad (9)$$




Figure 2 Loss landscape of Integrated Spatial Attack on CIFAR-10. (Left) A distant view of loss landscape w.r.t w before the optimization in Eq. 8. (Middle) A close view before the optimization shows a highly convex surface near the initialization point. (Right) The loss landscape around the maxima w^* after the optimization in Eq. 8.

where $w^{t+1} = [w_F^{t+1}, w_{RT}^{t+1}]^T$ is element-wisely clipped from \bar{w}^{t+1} by $\epsilon = [\epsilon_F, \epsilon_{RT}]^T$. From Figure 1, we can observe that our Integrated Spatial Attack can construct both local and global spatial transformations on images. Thus, it can simultaneously yield local pixel-wise permutations and global shape transformations.

Then, we visualize the loss surface under this Integrated Spatial Attack leveraging “filter normalization” [18] as illustrated in Figure 2. We strictly follow the implementation from [18] to achieve the desired visualization of the loss landscape of our integrated adversarial attacks for all the differentiable parameters w . Specifically, we view w_F and w_{RT} as two parameterized filters, which is analogous to the “filter normalization??” technique proposed by [18]. In the left part of Figure 2, we adjust the initialization of the variance of w , which then can provide a distant view of loss landscape before the optimization in Eq. 8. It exhibits a highly regular loss landscape, and its non-concavity w.r.t. only rotation and translation [5] has been tremendously improved. In the middle of Figure 2, we then provide a closer view of the loss landscape before the optimization. It shows a highly convex surface around the w to be optimized, facilitating the following optimization. In the right part of Figure 2, we also present the loss landscape around the maxima w^* after the optimization in Eq. 8 of our integrated spatial attack, exhibiting a highly concave surface as well. In summary, the highly non-concave loss landscape concerning only rotation and translation raised by [5] has been largely alleviated by considering both local and global spatial vulnerabilities. This integrated form smooths the optimization process, which guarantees the efficacy of our Integrated Spatial Attack.

2.4 Spatial Adversarial Training

As Eq. 9 incorporates local and global spatial robustness simultaneously, it is natural to leverage it to construct Spatial Adversarial Training, which we deploy in Experiment 4.4.

3 Relationship Between Sensitivity and Spatial Robustness

In this section, we will empirically investigate the relationships between different robustnesses and then explain them from the perspective of shape-based representation by leveraging a saliency map.

3.1 Relationships

We conduct rigorous experiments on MNIST, CIFAR-10, and Caltech-256 datasets to empirically examine the behavior of local and global spatial robustness as the sensitivity-based robustness increases. Specifically, after adversarially training multiple PGD (sensitivity-based) robust models with different numbers of PGD iterations, we further compute their test accuracy under Flow-based and RT-based spatial attacks via methods proposed in Section 2. The accuracy is computed on correctly classified test data for the model under consideration to mitigate the impact of the slightly different generalizations of these PGD-trained models. We fix both ϵ_F and ϵ_{RT} as 0.3 on MNIST, and choose ϵ_F and ϵ_{RT} as 0.3 and 1.0, respectively, on CIFAR-10 and Caltech-256. Then, we can control their strength of perturbations by adjusting the number of iterations in Flow-based and RT-based spatial attacks.

In Figure 3, the X-axis shows adversarially PGD-trained models with different numbers of PGD iterations, which can measure the different strengths of a model’s PGD (sensitivity-based) robustness. The Y-axis represents the computed test accuracy of the corresponding PGD-trained models under different

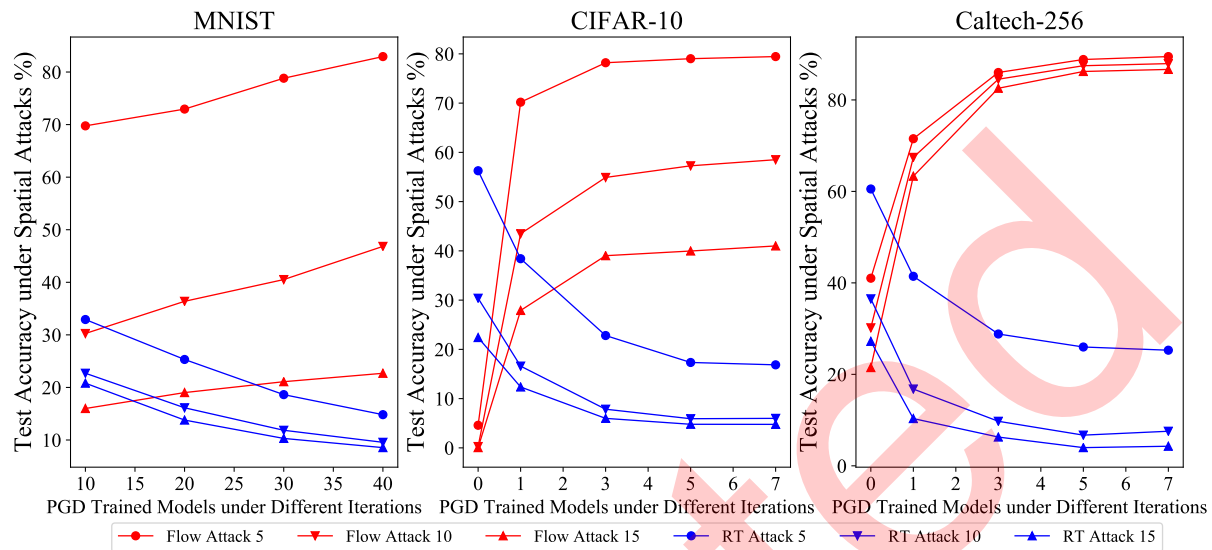


Figure 3 Relationships between sensitivity and two spatial robustness for three datasets. The X-axis represents adversarially PGD-trained models under different numbers of PGD iterations to measure the strength of sensitivity-based robustness, while the Y-axis represents the test accuracy under Flow Attack (red) and RT Attack (blue) with different iterations to measure the spatial robustness.

spatial attacks, and a high-level test accuracy reflects a model's high spatial robustness. It turns out that Flow-based spatial robustness (red lines) presents a steady ascending tendency across three datasets as the PGD sensitivity-based robustness increases, while the trend of RT-based spatial robustness (blue lines) fluctuates conversely. This result reveals that the sensitivity-based and RT-based spatial robustness is a trade-off relationship, consistent with the previous conclusion [12, 31]. However, this trade-off does not (even on the contrary) apply to the local spatial sensitivity, where sensitivity-based and Flow-based spatial robustness is positively correlated. We provide strong and intuitive evidence from the perspective of shape-biased representation below.

3.2 Explanation from the Shape-bias Representation

We show first with our brief conclusion: the sensitivity-based robustness corresponds to the sparse and shape-bias representation [26, 41], indicating that sensitivity-based robust models rely more on the *global shape* during prediction rather than the *local texture*. Nevertheless, the local and global spatial robustness are associated with different representation manners.

We visualize the saliency maps of naturally trained, PGD, Flow-based, and RT adversarially trained

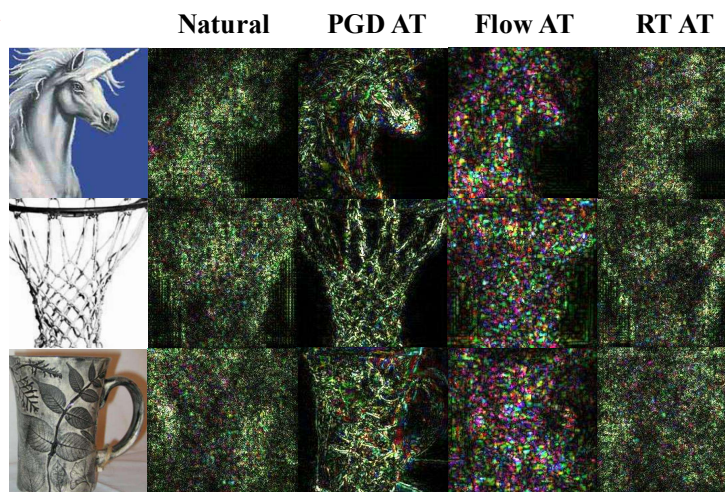


Figure 4 Saliency maps of four types of training models on some randomly selected images on Caltech-256.

models on some randomly selected images on Caltech-256, which are exhibited in Figure 4 to examine the shape-biased representation. Specifically, visualizing the saliency maps aims at assigning a sensitivity value, sometimes also called “attribution”, to show the sensitivity of the output to each pixel of an input image. Following [26, 41], we leverage SmoothGrad [27] to calculate the saliency map $S(x)$ of an image x , which alleviates the noises in the gradient by averaging over the gradient of n noisy copies of an input:

$$S(x) = \frac{1}{n} \sum_{i=1}^n \frac{\partial f_{\theta}^y(x_i)}{\partial x_i}, \tag{10}$$

where $x_i = x + q_i$, and q_i are noises drawn i.i.d from a Gaussian distribution $\mathcal{N}(0, \sigma^2)$. In our experiment, we set $n = 100$ and the noise level $\sigma/(x_{max} - x_{min}) = 0.1$.

Figure 4 shows that PGD-trained models tend to learn a *sparse and shape-biased* representation for all pixels of an image, while two types of spatially adversarially trained models suggest a converse representation. In particular, the representation from the Flow-based training model presents a *noisy and shape-biased* one as it places *extreme values*, although noisy, on pixels around the shape of objects, e.g., the edge between the horse and the background shown in Flow AT in Figure 4. On the contrary, RT-based models rely less on the shape of objects, and the saliency values tend to be *dense, smoothly* scattering around more pixels of an image.

We calculate the distance of saliency maps from different models across all test data on the Caltech-256 dataset and then compute their skewness in Figure 5. Specifically, we compute the pixel-wise distance between the saliency maps of the two models, and then we calculate the median of the skewness of the saliency map difference for all test data. Note that if two saliency maps have no statistical difference, then the difference in the values will follow a symmetric normal distribution with skewness 0. Negative skewness indicates that the original saliency map (representation) is sparse as compared to the model under consideration. We plot the tendency of skewness as the strength of some specific robustness increases in Figure 5. We summarize the observations into two conclusions:

1. Based on the first and fourth sub-pictures, both PGD and Flow-based robust models tend to learn a sparse and shape-biased representation compared with the natural model. However, the Flow-based trained model is less sparse (we call it noisy shape-biased) in comparison with the PGD-trained one.
2. In contrast, RT-based robust models tend to learn a dense representation. This is intuitive because the RT-trained model is expected to *memorize* broader pixel locations to cope with potential rotations and transformations in the test data.

Overall, the divergent representation (sparse vs. dense) between RT-based and sensitivity robustness verifies that the trade-off shown in Figure 3 is fundamental. More importantly, the positive correlation of sensitivity-based and local spatial robustness, shown in Figure 3, can also be explained by their similar shape-biased representation, although the latter tends to be noisy.

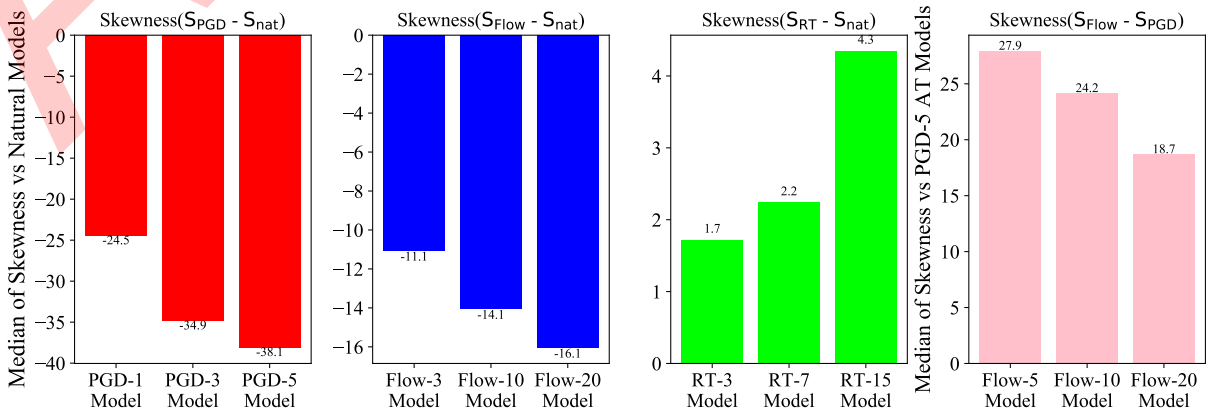


Figure 5 Median of skewness of saliency maps difference among robust models across all test data as compared with other models. The first three sub-pictures are compared with the naturally trained model, while the last one is compared with the PGD-trained model.

4 Pareto Adversarial Robustness

4.1 Motivation

Multi-objective Optimization. Given the insights garnered from our analysis of the relationships between natural accuracy and different kinds of adversarial robustness, a natural question that comes up is how to design a training strategy that can perfectly balance their mutual impacts, which mainly results from their different representation manners. In most cases, their relationships exhibit trade-offs, except for the positive correlation between sensitivity robustness and local spatial robustness. We use \mathcal{L}_{nat} , \mathcal{L}_{PGD} , $\mathcal{L}_{\text{Flow}}$ and \mathcal{L}_{RT} to represent the natural loss, the PGD adversarial loss, the Flow-based and the RT-based adversarial loss, respectively. We cast obtaining universal adversarial robustness as well as maintaining natural generalization ability as a multi-objective optimization problem [16], encompassing all of the aforementioned losses with a loss vector:

$$\min_{\theta} \mathcal{L}^{\theta} = (\mathcal{L}_0^{\theta}, \mathcal{L}_1^{\theta}, \mathcal{L}_2^{\theta}, \mathcal{L}_3^{\theta})^{\top}, \quad (11)$$

where $\mathcal{L}_0^{\theta}, \mathcal{L}_1^{\theta}, \mathcal{L}_2^{\theta}, \mathcal{L}_3^{\theta}$ represent $\mathcal{L}_{\text{nat}}, \mathcal{L}_{\text{PGD}}, \mathcal{L}_{\text{Flow}}, \mathcal{L}_{\text{RT}}$ respectively for simplicity, sharing the same model parameter θ . The multi-objective optimization is to optimize all loss functions simultaneously by exploiting the shared knowledge and structure, e.g., the representation.

Pareto Optimization. To harmonize these competing optimization objectives in the context of adversarial robustness, we introduce Pareto Optimization [14, 17, 19], which is successfully applied when optimal decisions need to be taken in the presence of trade-offs between multiple conflicting objectives. Pareto optimization endeavors to achieve Pareto Optimality, a balanced situation between all objectives, where none of the objective functions can be improved in value without degrading some of the other objective values. Mathematically, we have the following definitions [19, 42].

Pareto Dominance in Adversarial Robustness. Let θ^1, θ^2 be two parameters in the space Ω . θ^1 dominates θ^2 , i.e., $\theta^1 \prec \theta^2$, if and only if $\mathcal{L}_i^{\theta^1} \leq \mathcal{L}_i^{\theta^2}, \forall i \in \{0, 1, 2, 3\}$ and $\mathcal{L}_j^{\theta^1} < \mathcal{L}_j^{\theta^2}, \exists j \in \{0, 1, 2, 3\}$.

Pareto Optimality. θ^* is a Pareto optimal point, and \mathcal{L}^{θ^*} is a Pareto optimal objective vector if it does not exist $\hat{\theta} \in \Omega$ such that $\hat{\theta} \prec \theta^*$. The resulting Pareto front contains all Pareto optimal solutions.

Pareto Adversarial Robustness. Based on the insights presented above, a natural approach for incorporating Pareto criteria into multi-objective optimization in the context of adversarial training is to achieve universal adversarial robustness as well as maintain a desirable natural accuracy. The resulting Pareto Front contains all optimal, adversarially trained models for the given different constraints. The detailed formulation is presented later in Section 4.3.

4.2 Limitations of the Existing Strategies.

We denote $\mathcal{R}_{\text{adv}}(f; S_i) := \mathbb{E}_{(x,y) \sim \mathcal{D}} [\max_{r \in S_i} \mathcal{L}(f(x+r), y)]$ as the adversarial risk under perturbation sets $S_i, i = 1, \dots, m$. Our goal is to find f_{θ} that can achieve uniform risk minimization across all S_i as well




Figure 6 The difference between the model trained by the PGD method and Max AT with different parameter ϵ for the PGD attack in the PGD adversarial training.

as the minimal risk in the natural data. There are two common strategies to handle this issue.

1) Average adversarial training (Ave AT) [31]. $\mathcal{R}_{\text{ave}}(f; S) := \mathbb{E}_{(x,y) \sim \mathcal{D}} [\frac{1}{m} \sum_{i=1}^m \max_{r \in S_i} \mathcal{L}(f(x+r), y)]$, regards each adversarial robustness as having equal status. Intuitively, it may yield unsatisfactory solutions when the strength of different attacks mixed in the training are not balanced.

2) Max adversarial training (Max AT) [22,31], i.e., $\mathcal{R}_{\text{max}}(f; S) := \mathbb{E}_{(x,y) \sim \mathcal{D}} [\max_i \{ \max_{r \in S_i} \mathcal{L}(f(x+r), y) \}]$ tries to optimize over the max loss from the largest perturbations.

Overfitting issue of Max AT. Intuitively, Max AT may overfit to one specific type of adversarial robustness if its adversarial attack used for training is too strong. In Figure 6, we plot the difference in robust accuracy between Max AT and single PGD adversarial training. It turns out that as the strength of PGD attack ϵ used in Max AT increases, the difference among the three kinds of robust accuracy between Max AT and a single PGD AT tends to vanish. This indicates that the comprehensive robustness of Max AT degenerates to a single PGD adversarial training because the PGD loss tends to dominate as the strength of the PGD attack increases.

Overfitting issue of Ave AT based on its Relationship with Max AT. We consider the generalization issue based on different risks and then set the risk in Max AT and Ave AT as $\mathcal{R}_{\text{max}} = \max_i \mathcal{R}(f, S_i) = \max_i \mathcal{R}^{S_i}$ and $\mathcal{R}_{\text{ave}} = \frac{1}{m} \sum_{i=1}^m \mathcal{R}(f_{\theta}, S_i) = \frac{1}{m} \sum_{i=1}^m \mathcal{R}^{S_i}$. Proposition 2 informs that Max AT is closely associated with some form of Ave AT. This indicates that Max AT is likely to perform similarly to the specific form of Ave AT, which also suffers from unsatisfactory solutions when the strength of different attacks mixed in training is imbalanced. We verified this claim in Table 2 under a stronger PGD attack in Section 4.3.

Proposition 2. Given KKT differentiability and qualification conditions, $\exists \lambda_i \geq 0$, such that the risk minimizer in Max AT, i.e., $\mathcal{R}_{\text{max}}^*$ is a first-order stationary point of $\sum_{S_i \in \mathcal{S}} \lambda_i \mathcal{R}^{S_i}$ regardless of the relationship of S_i .

Remark. We point out that both Ave AT and Max AT may suffer from the robustness overfitting issue and thus fail in certain scenarios. However, a clever combination choice among all involved adversarial losses has the potential to alleviate the overfitting issues, thus outperforming both Max AT and Ave AT in terms of universal robustness. Motivated by this, we propose Pareto Adversarial Training in the next section, which will provide strong empirical evidence to support this intuition.

4.3 Pareto Adversarial Training

We apply linear scalarization to solve the multi-objective optimization, which is the most commonly used approach. We denote $\alpha = (\alpha_0, \alpha_1, \alpha_2, \alpha_3)$ as the combination coefficients for various losses. Thus, the objective function is $\min_{\theta} \sum_{i=0}^3 \mathbb{E}_x [\alpha_i^*(\theta) \mathcal{L}_i^{\theta}]$. Further, within the context of Pareto optimality, our goal is to find optimal combinations α between natural accuracy, sensitivity-based, and spatial robustness to perfectly balance their mutual impacts during the whole training process. Furthermore, we train a model f_{θ} under the optimal combinations α^* of different losses, and the computation of α^* in training is also associated with different losses determined by model parameters θ . This implies a bilevel optimization problem with θ as the upper-level variable and α as the lower-level variable. In the construction of low-level optimization regarding α , we apply a two-moment objective function concerning all losses. We name this bi-level optimization as *Pareto Adversarial Training*, which is formulated as:

$$\begin{aligned} &\min_{\theta} \sum_{i=0}^3 \mathbb{E}_x [\alpha_i^*(\theta) \mathcal{L}_i^{\theta}], \\ &\text{s.t. } \alpha^* = \arg \min_{\alpha} \sum_{i=0}^3 \sum_{j=0}^3 \mathbb{E}_x [(\alpha_i \mathcal{L}_i^{\theta} - \alpha_j \mathcal{L}_j^{\theta})^2], r = \sum_{i=1}^3 \alpha_i \mathbb{E}_x [\mathcal{L}_i^{\theta}], \sum_{i=0}^3 \alpha_i = 1, \alpha_i \geq 0, \forall i = 0, 1, 2, 3, \end{aligned} \tag{12}$$

where r indicates the expectation of one-moment over all robust losses, i.e., spatial and sensitivity-based losses, which reflects the strength of comprehensive robustness we require after solving this quadratic lower-level optimization regarding α . In particular, given the model parameter θ in each training step, the larger r we require will push the resulting $\alpha_i, i = 1, 2, 3$ larger, thus increasing the weight of the robust losses rather than the natural loss to pursue more robustness.

Algorithm 1 Bi-level Optimization in Pareto Adversarial Training.

Input: Training data $(\mathcal{X}, \mathcal{Y})$. Batch size M and adjustable hyper-parameter r . Initialization of α as $[1/4, 1/4, 1/4, 1/4]$.

Output: Classifier f_θ .

```

1: repeat
2:   Sample  $\{\mathbf{y}_1, \dots, \mathbf{y}_M\}$  and  $\{\mathbf{x}_1, \dots, \mathbf{x}_M\}$  from all training data.
3:   / * Step 1: Compute loss in Eq. 12 * /
4:   Compute natural loss  $\mathcal{L}_{\text{nat}}$ , and adversarial loss  $\mathcal{L}_{\text{PGD}}, \mathcal{L}_{\text{Flow}}, \mathcal{L}_{\text{RT}}$  based on natural cross entropy
   loss, PGD loss and Eq. 2 and 6, respectively.
5:   / * Step 2: Upper-level Optimization over  $\theta$  * /
6:   Given the current  $\alpha$ , update  $f_\theta$  by descending its stochastic gradient of:

       
$$\frac{1}{M} \sum_{i=1}^M \mathcal{L}^{\text{CE}}(f_\theta(\mathbf{x}_i), \mathbf{y}_i) = \frac{1}{M} \sum_{i=1}^M \alpha_0 \mathcal{L}_{\text{nat}}(f_\theta(\mathbf{x}_i), \mathbf{y}_i) + \alpha_1 \mathcal{L}_{\text{PGD}}(f_\theta(\mathbf{x}_i), \mathbf{y}_i) \\ + \alpha_2 \mathcal{L}_{\text{Flow}}(f_\theta(\mathbf{x}_i), \mathbf{y}_i) + \alpha_3 \mathcal{L}_{\text{RT}}(f_\theta(\mathbf{x}_i), \mathbf{y}_i)$$


7:   / * Step 3: Lower-level Optimization over  $\alpha$  * /
8:   Compute  $\hat{\mu}$  and  $\hat{\Sigma}$  by sliding window technique in Eq. 13.
9:   Evaluate  $P$  in the quadratic form shown in Eq. 13.
10:  Solve Eq. 13 via CVXOPT tool to obtain the  $\alpha$ .
11: until Convergence

```

Two-Moment Objective Function. The two-moment form is a common practice in Pareto optimization. For example, in the financial portfolio theory, the mean-variance optimization is normally leveraged to compute the *Pareto Efficient Front*, where the risk of the asset portfolio, measured by their variances, is minimized to balance the different correlations of these assets given an expected return from the investor. Similarly, the square loss of the difference between each loss pair in Eq. 12 measures their mutual impacts. For instance, a decrease in \mathcal{L}_{PGD} tends to increase \mathcal{L}_{RT} as they have a fundamental trade-off relationship. We hope to mitigate all these mutual impacts, measured by the weighted quadratic differences, among all losses given an expected robustness level of r . In the implementation, as we regard all losses as random variables with their stochasticity arising from the mini-batch sampling from data, we leverage the sliding windows technique to compute their expectations. Our bi-level optimization within a batch is (1) θ : update parameters θ via SGD and (2) α : solve α via quadratic programming. Denote the random variables $\mathcal{L}_0, \mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$ with mean vector μ and covariance matrix Σ . We transform our lower-level optimization regarding α as the following standard quadratic form:

$$\min_{\theta, \alpha} \alpha^T P \alpha \quad \text{s.t.} \quad \begin{bmatrix} 0 & \mu_1 & \mu_2 & \mu_3 \\ 1 & 1 & 1 & 1 \end{bmatrix} \alpha = \begin{bmatrix} r \\ 1 \end{bmatrix}, -\alpha \leq \mathbf{0}, \quad (13)$$

where $P = 8(\text{diag}(\Sigma) + \text{diag}(\mu\mu^T)) - 2(\Sigma + \mu\mu^T)$. We utilize the *CVXOPT* tool to solve this quadratic optimization within each mini-batch training. *CVXOPT* is probably the most popular free software package for convex optimization based on the Python programming language that can solve quadratic programming effectively. We also provide proof of the quadratic formulation in Appendix D.

A detailed algorithm description is given in Algorithm 1. In the lower-level procedure of Pareto adversarial training, we solve the quadratic optimization regarding α given θ in each training step to obtain the optimal combinations among natural loss, sensitivity-based, and spatial adversarial loss. Then in the upper-level optimization, we leverage our familiar SGD method to update θ based on α^* calculated from the lower-level problem. Note that the computation complexity of our method is similar to Ave AT and Max AT, which are still competitive in computation.

4.4 Approximated Pareto Front

By adjusting the upper bound of the expected adversarial robustness loss r , we can evenly generate Pareto optimal solutions where the obtained models will have different levels of robustness under optimal combinations. The set of all Pareto optimal solutions then forms the *Pareto front*. Rigorously, it is almost impossible to attain all Pareto optimal solutions for a general continuous multi-objective optimization problem unless a closed-form solution exists for each r . Alternatively, we leverage the limited solutions

Dataset	Robustness Score (%)	Natural Model	PGD AT	Spatial AT	Max AT	Ave AT	Pareto AT ($r = 2.2$)
MNIST	Sensitivity-based Robustness	29.40	98.42	0.24	65.16	92.70	88.06
	Local Spatial Robustness	14.36	38.23	27.59	53.02	48.51	58.70
	Global Spatial Robustness	16.70	12.77	78.76	51.47	88.76	90.40
	Universal Robustness	0.0	88.97	46.14	109.19	169.52	176.71
Dataset	Robustness Score (%)	Natural Model	PGD AT	Spatial AT	Max AT	Ave AT	Pareto AT ($r = 4.0$)
CIFAR-10	Sensitivity-based Robustness	0.82	70.24	12.11	52.04	49.68	51.65
	Local Spatial Robustness	9.34	72.29	83.63	85.31	79.88	80.38
	Global Spatial Robustness	57.28	18.94	40.96	40.06	64.98	66.36
	Universal Robustness	0.0	94.04	69.26	109.98	127.11	130.96
Dataset	Robustness Score (%)	Natural Model	PGD AT	Spatial AT	Max AT	Ave AT	Pareto AT ($r = 6.0$)
Caltech-256	Sensitivity-based Robustness	4.74	82.43	6.94	59.81	71.60	76.52
	Local Spatial Robustness	34.59	87.96	88.75	65.89	86.67	87.39
	Global Spatial Robustness	49.73	21.71	65.04	64.64	53.68	50.00
	Universal Robustness	0.0	103.05	71.66	101.28	122.89	124.85

Table 1 Robustness Score of each type of adversarial robustness on MNIST, CIFAR-10, and Caltech-256. Each type of robustness (%) is the average test accuracy under different strengths of perturbations. We choose the universal robustness of the Natural Model as the baseline and set it as 0. We use the difference in average test accuracy between other models and the Natural Model and then sum them as the universal robustness.

obtained by solving a series of multi-objective optimization problems for various r to approximate the Pareto front.

Thus, we train deep neural networks under different adversarial training strategies, i.e., PGD Adversarial Training (PGD AT), Spatial Adversarial Training (Spatial AT) proposed in Section 2.4, Max AT, Ave AT, and Pareto Adversarial Training (Pareto AT) under different r , in which we apply a proper iteration. Then we evaluate their test accuracy under PGD, Flow-based, and RT attacks under different perturbation strengths. Next, we average the test accuracies for each type of attack, and the result is a quantitative measure of the specific robustness, called *Robustness Score*. To evaluate the universal robustness, we further compute the average of Robustness Scores for all kinds of robustness and use the increment over the naturally trained model as the metric called *Universal Robustness Score*. We report the Robustness Scores of all models on CIFAR-10 in Table 1, and the results on the other two datasets are similar. All implementation details are provided in Appendix E. It shows that Pareto AT ($r = 4.0$) has the best universal robustness score among all the models considered, although the highest specific robustness normally exists in the adversarial training model that only focuses on it.

Finally, we plot the universal robustness scores and the sacrificed clean accuracy of all methods across three datasets in Figure 7, where multiple Pareto AT models (red points) are trained under different r . The Pareto criterion exhibited in Figure 7 provides an optimality principle, which enables Pareto Adversarial Training to achieve the best universal robustness among all the methods considered, given a certain tolerable level of sacrificed clean accuracy. By adjusting the expected universal robustness r in Pareto Adversarial Training, we can develop the set of Pareto optimal solutions, i.e., the Pareto front. It shows that all other methods are above our Pareto front and are less effective than our proposal.

Overfitting Issue of Ave AT. Note that although the perturbation strength adopted in Table 1 is




Figure 7 The Pareto front (red lines) between the universal robustness score and sacrificed clean accuracy on MNIST, CIFAR-10, and Caltech-256. The vertical axis is the decrease of the natural accuracy compared with the naturally trained model and has been under the log transformation along two directions.

Robustness Score (%)	Natural Model	Ave AT	Pareto AT ($r = 3.5$)	Pareto AT ($r = 3.7$)	Pareto AT ($r = 4.0$)	Pareto AT ($r = 4.2$)
Natural Accuracy	91.43	56.39	82.64	79.69	71.68	61.53
Sensitivity-based	0.82	64.11	53.73	58.70	63.28	65.19
Local Spatial	9.34	82.45	80.10	77.62	81.01	82.38
Global Spatial	57.28	51.36	66.57	67.56	59.69	52.04
Universal Robustness	0.0	197.92	200.39(+2.37)	203.88(+5.96)	203.98(+6.06)	199.61(+1.69)

Table 2 Robustness Score on CIFAR-10 with a **larger** step size 8/255 and ϵ as 16/255 in PGD perturbations used for both Ave AT and Pareto AT across different r .

mild, we need to point out that *the superiority of Pareto AT over Ave AT can be higher if the overfitting issue is severe*. We demonstrate this claim in Table 2, where we apply a **stronger** PGD perturbation used in AT. Finally, we find that Ave AT overfits sensitivity robustness more severely, achieving much less universal robustness and sacrificing more clean accuracy than Pareto AT. Pareto Adversarial Training can mitigate the overfitting issue regarding an overly strong perturbation in AT because Pareto AT can automatically adjust the weights α while training, which is the key advantage of Pareto AT over Ave AT. **Sensitivity Analysis.** Comparing universal adversarial robustness between Table 1 and Table 2, it can be seen that Pareto AT achieves more consistent universal adversarial robustness. In addition to this sensitivity analysis in terms of perturbation sizes, we also investigate the variation of universal adversarial robustness by changing the expected adversarial robustness loss r . Results are provided in Table 2. It suggests that Pareto AT with a mild r can achieve the best universal robustness score, while Pareto AT with an excessively large or small r may not have sufficient universal robustness. Moreover, Pareto Front in Figure 7 also serves as the sensitivity analysis results in terms of different r .

Overall, we conclude that Pareto adversarial training perfectly balances the mutual impacts of sensitivity-based robustness and spatial robustness under the Pareto criterion.

5 Discussion and Conclusion

The principal purpose of our work is to design a novel approach to achieve universal adversarial robustness. We first analyze the two main branches of spatial robustness and then integrate them into one attack and adversarial training design. Furthermore, we investigate the comprehensive relationships between sensitivity-based and two distinct spatial robustnesses from the perspective of representation. Based on the understanding of the mutual impacts of different kinds of adversarial robustness, we introduce the Pareto criterion into the adversarial training framework to develop Pareto Adversarial Training. The resulting Pareto front provides optimal solutions over existing baselines, given the universal robustness level we hope to attain. In the future, we hope to apply Pareto analysis to more general Out-of-Distribution generalization settings.

Acknowledgements Z. Lin was supported by the major key project of PCL (No. PCL2021A12) and the NSF China (No. 62276004).

References

- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.
- Zhaohui Che, Ali Borji, Guangtao Zhai, Suiyi Ling, Jing Li, Xiongkuo Min, Guodong Guo, and Patrick Le Callet. Smgea: A new ensemble adversarial attack powered by long-term gradient memories. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- Gavin Weiguang Ding, Yash Sharma, Kry Yik Chau Lui, and Ruitong Huang. Max-margin adversarial (mma) training: Direct input space margin maximization through adversarial training. *International Conference on Learning Representations, ICLR 2020*, 2018.
- Logan Engstrom, Brandon Tran, Dimitris Tsipras, Ludwig Schmidt, and Aleksander Madry. Exploring the landscape of spatial robustness. In *International Conference on Machine Learning*, pages 1802–1811, 2019.
- Logan Engstrom, Dimitris Tsipras, Ludwig Schmidt, and Aleksander Madry. A rotation and a translation suffice: Fooling cnns with simple transformations. *arXiv preprint arXiv:1712.02779*, 1(2):3, 2017.
- Yarin Gal. Uncertainty in deep learning. *University of Cambridge*, 1(3), 2016.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *International Conference on Learning Representations*, 2014.
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. *arXiv preprint arXiv:2006.16241*, 2020.

10 Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *International Conference on Learning Representations*, 2019.

11 Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015.

12 Sandesh Kamath, Amit Deshpande, and KV Subrahmanyam. Invariance vs. robustness of neural networks. *arXiv preprint arXiv:2002.11318*, 2020.

13 Il Yong Kim and OL De Weck. Adaptive weighted sum method for multiobjective optimization: a new method for pareto front generation. *Structural and multidisciplinary optimization*, 31(2):105–116, 2006.

14 Il Yong Kim and Oliver L De Weck. Adaptive weighted-sum method for bi-objective optimization: Pareto front generation. *Structural and multidisciplinary optimization*, 29(2):149–158, 2005.

15 David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). *arXiv preprint arXiv:2003.00688*, 2020.

16 Man-Fai Leung and Jun Wang. A collaborative neurodynamic approach to multiobjective optimization. *IEEE transactions on neural networks and learning systems*, 29(11):5738–5748, 2018.

17 Cong Li, Michael Georgiopoulos, and Georgios C Anagnostopoulos. Pareto-path multitask multiple kernel learning. *IEEE transactions on neural networks and learning systems*, 26(1):51–61, 2014.

18 Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *Advances in Neural Information Processing Systems*, pages 6389–6399, 2018.

19 Xi Lin, Hui-Ling Zhen, Zhenhua Li, Qing-Fu Zhang, and Sam Kwong. Pareto multi-task learning. In *Advances in Neural Information Processing Systems*, pages 12060–12070, 2019.

20 Qi Liu and Wujie Wen. Model compression hardens deep neural networks: A new perspective to prevent adversarial attacks. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.

21 Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *International Conference on Learning Representations, ICLR 2018*, 2017.

22 Pratyush Maini, Eric Wong, and J Zico Kolter. Adversarial robustness against the union of multiple perturbation models. *International Conference on Machine Learning*, 2019.

23 Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.

24 Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John Duchi, and Percy Liang. Understanding and mitigating the tradeoff between robustness and accuracy. *International Conference on Machine Learning*, 2020.

25 Mahmood Sharif, Lujo Bauer, and Michael K Reiter. On the suitability of lp-norms for creating and preventing adversarial examples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1605–1613, 2018.

26 Baifeng Shi, Dinghui Zhang, Qi Dai, Zhanxing Zhu, Yadong Mu, and Jingdong Wang. Informative dropout for robust representation learning: A shape-bias perspective. *International Conference on Machine Learning*, 2020.

27 Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017.

28 Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

29 Richard Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.

30 Florian Tramèr, Jens Behrmann, Nicholas Carlini, Nicolas Papernot, and Jörn-Henrik Jacobsen. Fundamental tradeoffs between invariance and sensitivity to adversarial perturbations. *ICML*, 2020.

31 Florian Tramèr and Dan Boneh. Adversarial training and robustness for multiple perturbations. In *Advances in Neural Information Processing Systems*, pages 5866–5876, 2019.

32 Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. *International Conference on Learning Representations, ICLR 2019*, 2018.

33 Vladimir N Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity*, pages 11–30. Springer, 2015.

34 Haotao Wang, Tianlong Chen, Shupeng Gui, Ting-Kuei Hu, Ji Liu, and Zhangyang Wang. Once-for-all adversarial training: In-situ tradeoff between robustness and accuracy for free. *Advances in Neural Information Processing Systems*, 2020.

35 Chaowei Xiao, Jun-Yan Zhu, Bo Li, Warren He, Mingyan Liu, and Dawn Song. Spatially transformed adversarial examples. *International Conference on Learning Representations, ICLR 2018*, 2018.

36 Nanyang Ye, Qianxiao Li, Xiao-Yun Zhou, and Zhanxing Zhu. An annealing mechanism for adversarial training acceleration. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.

37 Dong Yin, Raphael Gontijo Lopes, Jon Shlens, Ekin Dogus Cubuk, and Justin Gilmer. A fourier perspective on model robustness in computer vision. In *Advances in Neural Information Processing Systems*, pages 13276–13286, 2019.

38 Milan Zeleny. *Multiple criteria decision making Kyoto 1975*, volume 123. Springer Science & Business Media, 2012.

39 Haichao Zhang and Jianyu Wang. Joint adversarial training: Incorporating both spatial and pixel attacks. *arXiv preprint arXiv:1907.10737*, 2019.

40 Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P Xing, Laurent El Ghaoui, and Michael I Jordan. Theoretically principled trade-off between robustness and accuracy. *International Conference on Machine Learning*, 2019.

41 Tianyuan Zhang and Zhanxing Zhu. Interpreting adversarially trained convolutional neural networks. *International Conference on Machine Learning*, 2019.

42 Eckart Zitzler and Lothar Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE transactions on Evolutionary Computation*, 3(4):257–271, 1999.

Appendix A Visualization of Various Attacks

To better present the visual effect of various kinds of adversarial attacks, we provide high-resolution results on Caltech-256 in Figure A1. It turns out that Flow-based attacks focus on local spatial vulnerability that mainly blurs pixels in some local regions, while RT attacks cause a shape-based global spatial transformation. More importantly, our integrated spatial attacks are more comprehensive in the sense of spatial robustness, combining both local and local spatial sensitivity.




Figure A1 High-resolution images on Caltech-256.

Appendix B Proof of Proposition 1

Proof. Firstly, we have the following equations according to the definitions of the loss function:

$$\begin{aligned}\mathcal{L}_{\theta}^{\text{CE}}(x_{w_F}, y) &= \log \sum_{i=1}^K \exp(f_{\theta}^i(x_{w_F})) - f_{\theta}^y(x_{w_F}) \\ \mathcal{L}_{\theta}^S(x_{w_F}, y) &= \log \sum_{i \neq y} \exp(f_{\theta}^i(x_{w_F})) - f_{\theta}^y(x_{w_F})\end{aligned}\quad (\text{B1})$$

Then, we compute their gradients for the flow vector x_{w_F} . The gradient of $\mathcal{L}_{\theta}^{\text{CE}}(x_{w_F}, y)$ is shown as follows:

$$\begin{aligned}\nabla_{w_F} \mathcal{L}_{\theta}^{\text{CE}}(x_{w_F}, y) &= \frac{\sum_{i=1}^K \exp(f_{\theta}^i(x_{w_F})) \cdot \nabla_{x_{w_F}} f_{\theta}^i(x_{w_F}) \cdot \nabla_{w_F} x_{w_F}}{\sum_{i=1}^K \exp(f_{\theta}^i(x_{w_F}))} - \nabla_{x_{w_F}} f_{\theta}^y(x_{w_F}) \cdot \nabla_{w_F} x_{w_F} \\ &= \frac{\sum_{i=1}^K \exp(f_{\theta}^i(x_{w_F})) \nabla_{w_F} x_{w_F} (\nabla_{x_{w_F}} f_{\theta}^i(x_{w_F}) - \nabla_{x_{w_F}} f_{\theta}^y(x_{w_F}))}{\sum_{i=1}^K \exp(f_{\theta}^i(x_{w_F}))}\end{aligned}\quad (\text{B2})$$

Similarly, the gradient of $\mathcal{L}_{\theta}^S(x_{w_F}, y)$ is:

$$\nabla_{w_F} \mathcal{L}_{\theta}^S(x_{w_F}, y) = \frac{1}{\sum_{i \neq y} \exp(f_{\theta}^i(x_{w_F}))} \cdot \left(\sum_{i \neq y} \exp(f_{\theta}^i(x_{w_F})) \nabla_{w_F} x_{w_F} (\nabla_{x_{w_F}} f_{\theta}^i(x_{w_F}) - \nabla_{x_{w_F}} f_{\theta}^y(x_{w_F})) \right) \quad (\text{B3})$$

Then we take the multiplication of $\nabla_{w_F} \mathcal{L}_{\theta}^S(x_{w_F}, y)$ by a term $\frac{\sum_{i \neq y} \exp(f_{\theta}^i(x_{w_F}))}{\sum_{i=1}^K \exp(f_{\theta}^i(x_{w_F}))}$, finally we attain:

$$\begin{aligned}\nabla_{w_F} \mathcal{L}_{\theta}^S(x_{w_F}, y) \cdot \frac{\sum_{i \neq y} \exp(f_{\theta}^i(x_{w_F}))}{\sum_{i=1}^K \exp(f_{\theta}^i(x_{w_F}))} &= \frac{\sum_{i \neq y} \exp(f_{\theta}^i(x_{w_F})) \nabla_{w_F} x_{w_F} (\nabla_{x_{w_F}} f_{\theta}^i(x_{w_F}) - \nabla_{x_{w_F}} f_{\theta}^y(x_{w_F})) + 0}{\sum_{i=1}^K \exp(f_{\theta}^i(x_{w_F}))} \\ &= \frac{\sum_{i=1}^K \exp(f_{\theta}^i(x_{w_F})) \nabla_{w_F} x_{w_F} (\nabla_{x_{w_F}} f_{\theta}^i(x_{w_F}) - \nabla_{x_{w_F}} f_{\theta}^y(x_{w_F}))}{\sum_{i=1}^K \exp(f_{\theta}^i(x_{w_F}))} \\ &= \nabla_{w_F} \mathcal{L}_{\theta}^{\text{CE}}(x_{w_F}, y)\end{aligned}\quad (\text{B4})$$

Finally, we denote $\frac{\sum_{i \neq y} \exp(f_{\theta}^i(x_{w_F}))}{\sum_{i=1}^K \exp(f_{\theta}^i(x_{w_F}))}$ as $r(x_{w_F}, y)$.

Appendix C Proof of Proposition 2

Proof. Let f as the minimizer, e.g., the neural networks after the optimization:

$$\begin{aligned} f^* &\in \min_f \max_i \mathcal{R}(f, S_i) \\ M^* &= \max_i \mathcal{R}(f, S_i) \end{aligned} \quad (C1)$$

Then the optimization can be equivalent to a constrained version:

$$\begin{aligned} \min_{f, M} \quad & M \\ \text{s.t.} \quad & \mathcal{R}(f, S_i) \leq M \text{ for all } S_i \in \mathcal{S} \end{aligned} \quad (C2)$$

with Lagrangian $L(f, M, \lambda) = M + \sum_{S_i \in \mathcal{S}} \lambda_i (\mathcal{R}(f, S_i) - M)$. If this optimization problem satisfies KKT condition, then $\exists \lambda_i \geq 0$ with $\nabla_f L(f^*, M^*, \lambda) = 0$ such that

$$\nabla_f|_{f=f^*} \sum_{S_i \in \mathcal{S}} \lambda_i \mathcal{R}(f, S_i) = 0$$

Remark We point out that our conclusion is made under the assumption that the KKT condition holds and the stationary point of f regarding the Lagrangian function can be attained, which normally requires the convexity condition. However, under these assumptions, we can still establish the close correlation between Max AT and Ave AT, indicating they are likely to perform similarly in many cases.

Appendix D Optimization analysis on the Pareto Adversarial Training and Algorithm

We provide the proof of P in the following:

Proof.

$$\begin{aligned} \sum_{i=0}^3 \sum_{j=0}^3 \mathbb{E}(\alpha_i \mathcal{L}_i - \alpha_j \mathcal{L}_j)^2 &= \sum_{i=0}^3 \sum_{j=0}^3 \mathbb{E}((\alpha_i \mathcal{L}_i - \mathbb{E}(\alpha_i \mathcal{L}_i)) - (\alpha_j \mathcal{L}_j - \mathbb{E}(\alpha_j \mathcal{L}_j)) + (\mathbb{E}(\alpha_i \mathcal{L}_i) - \mathbb{E}(\alpha_j \mathcal{L}_j)))^2 \\ &= \sum_{i=0}^3 \sum_{j=0}^3 \mathbb{E}((\alpha_i \mathcal{L}_i - \mathbb{E}(\alpha_i \mathcal{L}_i)) - (\alpha_j \mathcal{L}_j - \mathbb{E}(\alpha_j \mathcal{L}_j)))^2 + (\mathbb{E}(\alpha_i \mathcal{L}_i) - \mathbb{E}(\alpha_j \mathcal{L}_j))^2 + 0 \\ &= \sum_{i=0}^3 \sum_{j=0}^3 (\alpha_i^2 \sigma_{ii} + \alpha_j^2 \sigma_{jj} - 2\alpha_i \alpha_j \sigma_{ij}) + (\alpha_i^2 \mu_i^2 + \alpha_j^2 \mu_j^2 - 2\alpha_i \alpha_j \mu_i \mu_j) \\ &= 8\alpha^T \text{diag}(\Sigma) \alpha - 2\alpha^T \Sigma \alpha + 8\alpha^T \text{diag}(\mu \mu^T) \alpha - 2\alpha^T (\mu \mu^T) \alpha \\ &= \alpha^T (8(\text{diag}(\Sigma) + \text{diag}(\mu \mu^T)) - 2(\Sigma + \mu \mu^T)) \alpha \end{aligned} \quad (D1)$$

Appendix E Implementation

Implementation Details. For MNIST comparison, we train the Simple CNN in [40] on MNIST for 100 epochs. As for the CIFAR-10 dataset, we choose the widely used Pre-Act ResNet-18 with grouped normalization and trained the network for 76 epochs. The other details of our implementation on MNIST and CIFAR-10 are based on [40], while the implementation on Caltech-256 has to refer to [41] with 10 epochs to finetune a pre-trained ResNet-18.

- **PGD Attack.** We apply the widely accepted setting on these three datasets. We set step size as 0.01, ϵ as 0.3 on MNIST while the step size is 0.007 and ϵ is 0.031 on both CIFAR-10 and Caltech-256 datasets. To evaluate the different levels of robustness, we evaluate PGD attack under 10, 20, 30, 40 iterations on MNIST and 5, 10, 15, 20 iterations on CIFAR-10 and Caltech-256 datasets.
- **Flow-based and RT Attacks.** On MNIST, we set step size α_F and α_{RT} as 0.01 and 0.1, and choose $\epsilon_F, \epsilon_{RT}$ as 0.3. We select 5, 10, 15, 20 as the attack iterations for the evaluation of both two attacks. On CIFAR-10, we set step size α_F as $1e-3$ and α_{RT} as 0.05, and choose $\epsilon_F, \epsilon_{RT}$ as 0.3, 1.0. We select 3, 5, 10, 15 as the attack iterations for the evaluation of both two attacks. On Caltech-256, we set step size α_F as $1e-5$ and α_{RT} as 0.1, and choose $\epsilon_F, \epsilon_{RT}$ as 0.3 and 1.0 for the two attacks, respectively. We select 3, 5, 10, 15 as the attack iterations for the evaluation of both two attacks.
- **PGD Adversarial Training.** We choose PGD iterations as 30, 3, and 5 in the PGD adversarial training on MNIST, CIFAR-10, and Caltech-256, respectively. The adversarial attack strength is the same as PGD attacks for each dataset, respectively.
- **Spatial Adversarial Training.** Our integrated spatial adversarial training is based on our proposed integrated spatial attacks that unify both Flow-based and RT-based attacks. We set the iterations as 20, 5, 10 and on MNIST, CIFAR-10 and Caltech-256, respectively. Other hyper-parameters are the same as those in their corresponding attacks.
- **Pareto Adversarial Training.** The parameter r is the measure of comprehensive adversarial robustness. We select a sequence of r to train multiple Pareto Adversarial training models. Particularly, on MNIST, we choose r in $[0.2, 0.5, 0.8, 1.0, 1.2, 1.5, 1.8, 2.0, 2.2]$, and r in $[0.5, 1.0, 1.25, 1.5, 2.25, 3.0, 3.5, 4.0]$ on CIFAR-10 and Caltech-256. Other parameters follow the corresponding methods above, respectively.