

Enhancing Adversarial Robustness via Test-time Transformation Ensembling

Juan C. Pérez^{1,2}, Motasem Alfarra¹, Guillaume Jeanneret², Laura Rueda²,
 Ali Thabet¹, Bernard Ghanem¹, and Pablo Arbeláez²

¹King Abdullah University of Science and Technology (KAUST),

²Center for Research and Formation in Artificial Intelligence, Universidad de los Andes,

¹{juan.perezsantamaria,motasem.alfarra,ali.thabet,bernard.ghanem}@kaust.edu.sa

²{g.jeanneret10,l.ruedag,pa.arbelaez}@uniandes.edu.co;

Abstract

Deep learning models are prone to being fooled by imperceptible perturbations known as adversarial attacks. In this work, we study how equipping models with Test-time Transformation Ensembling (TTE) can work as a reliable defense against such attacks. While transforming the input data, both at train and test times, is known to enhance model performance, its effects on adversarial robustness have not been studied. Here, we present a comprehensive empirical study of the impact of TTE, in the form of widely-used image transforms, on adversarial robustness. We show that TTE consistently improves model robustness against a variety of powerful attacks without any need for re-training, and that this improvement comes at virtually no trade-off with accuracy on clean samples. Finally, we show that the benefits of TTE transfer even to the certified robustness domain, in which TTE provides sizable and consistent improvements.

1. Introduction

The onset of deep learning techniques has revolutionized several fields such as Computer Vision [28], Natural Language Processing [34], and Reinforcement Learning [35]. In the realm of computer vision, deep learning based methods have even surpassed human-level performance on challenging datasets [21]. Despite the success of deep learning-based systems, researchers have noticed severe brittleness in their output: while remarkably accurate, they are extremely sensitive to imperceptible perturbations, now known as adversarial attacks [51].

The discovery of this adversarial vulnerability, accordingly, has led to a large number of works addressing security concerns by proposing methods to defend models against attacks [7, 19, 32, 18]. Beyond such security concerns, adversarial vulnerability also sheds doubt on the nature of the impressive performance that computer vision systems achieve. While some of these systems may have been inspired by the human visual system, their lack of robustness

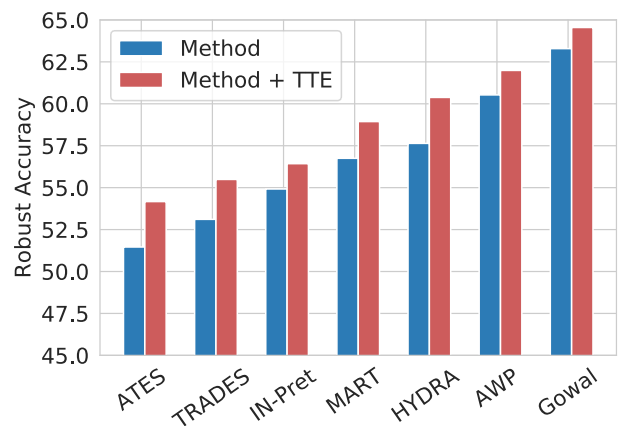


Figure 1. **Test-time Transformation Ensembling (TTE) enhances adversarial robustness.** Introducing TTE increases the adversarial robustness of several state-of-the-art defenses on CIFAR10.

could suggest that their inner workings strongly differ from those of humans [56].

We highlight that current recognition methods perform predictions on single and static images. While this experimental setup is practical, it strongly differs from the way in which humans recognize objects in the real world. We note that this mismatch in how models perform predictions may be a factor contributing to the adversarial vulnerability: studies in developmental psychology have noted how toddlers jointly exploit object permanence [41, 5] and multiple views of objects to develop efficient, accurate and *robust* visual systems [49, 33, 24, 40, 4]. That is, while humans may perform recognition by ensembling multiple views of the intended object, machine learning-based systems traditionally focus on a single view. Inspired by these observations, we argue that ensembling predictions of transformed versions of an image can be cast as a simple and coarse simulation of how humans recognize objects in the real world.

Indeed, studies have shown that harnessing transform-

based augmentation of data, both at train and test time, can provide consistent performance boosts in tasks such as detection [30] and recognition [28, 47, 22]. Other works have shown that randomization-based transforms and noise injection can be exploited to improve adversarial robustness [42, 31], providing evidence that randomization is useful for defending against adversaries.

In this work, we study the effect that Test-time Transformation Ensembling (TTE) of model predictions on *deterministically*-transformed versions of an image has on adversarial robustness. Specifically, we explore ensembling predictions over transformed versions of an image by applying two semantics-preserving transforms customarily used in computer vision (*crops* and *flips*). We present a comprehensive empirical study assessing the impact of introducing TTE against state-of-the-art adversaries. Our experiments show that equipping deep networks with TTE provides consistent improvements in adversarial robustness across datasets and attacks, while coming at virtually no cost to accuracy on clean samples. TTE is modular and versatile, and we show how it improves robustness of state-of-the-art defenses across-the-board (Figure 1). We further validate the versatility of TTE by showing how it can also boost certified robustness.

Contributions: (1) We present Test-time Transformation Ensembling (TTE), where we apply two semantics-preserving transforms (*crops* and *flips*), to enhance network robustness against adversaries. Upgrading a model with TTE does not need re-training and it can be implemented with less than 30 lines of code. To ensure reproducibility, we will make our full implementation in PyTorch [39] publicly available ¹ (2) We show how TTE consistently provides significant improvements in adversarial robustness for top-performing methods on the *AutoAttack* benchmark [13], reaching improvements of 2.04% and 1.29% on CIFAR10 and CIFAR100, respectively (Figure 1). TTE’s benefits also extend to the large-scale ImageNet dataset [15], where we enhance the state-of-the-art defense, Feature Denoising [56], and conduct the first large-scale evaluation of adversarial robustness against *AutoAttack* on ImageNet. We find that, even in the ImageNet domain, TTE can provide boosts of 2.21%. (3) We find that the benefits of TTE are not confined to the *empirical* assessment of adversarial robustness: TTE can also boost *certified* defenses. Specifically, on CIFAR10, introducing TTE into Randomized Smoothing [11] and SmoothAdv [44] provides a boost of 10% and 7% in Average Certified Radius, respectively.

2. Related work

Data transforms. Classic transforms are behind most augmentation techniques for preventing overfitting of Deep

Neural Networks [28]. These transforms are now a *de facto* standard in training-routine implementations, and are of such importance that automated learning of augmentation policies has been studied [14]. While transforms are effective during training, a stream of works also showed their benefits at test time [22, 50, 47, 28, 25, 37]. For instance, Sermanet *et al.* [46] augment the input at test time by including multi-scale information. Moreover, [30, 59] enhanced detectors by ensembling the predictions on both the original images and their transformed versions. In this work, we observe that data transforms are related to the “self-generated variability” phenomenon observed in infants [49]. Inspired by this connection, we present a comprehensive empirical study of the effects of test-time transforms on adversarial robustness.

Adversarial Robustness. Robustness to adversarial noise is a pressing concern in the deep learning community [19, 16]. However, evaluating adversarial robustness has proven to be extremely difficult [6]. A large amount of works have proposed both white-box attacks [29, 32, 36] and black-box attacks [1, 2, 3, 17] to evaluate adversarial robustness. Recently, Croce *et al.* [13] proposed *AutoAttack*, an ensemble of four attacks with automatic hyper-parameter calibration that provides reliable assessments of adversarial robustness. From the side of defenses, a stream of works have aimed at developing models that are resistant to attacks. These works include adversarial training [19, 32], combining adversarial training with pre-training on external data sources [23, 10, 26, 8], and TRADES [61], among many others [55, 53, 45, 9]. In this work, we study how defenses can be equipped with TTE to boost adversarial robustness as measured via *AutoAttack*.

Certified Defenses. The aforementioned empirical assessment of adversarial robustness cannot provide guarantees about the inexistence of adversarial examples for a given classifier. This concern has incited interest in certified adversarial robustness, which aims at developing models whose predictions can be certified, *i.e.* predictions that are verifiably constant within some radius around any input [54, 43]. Here, we note that many approaches have been proposed for the purpose of certification [60, 11, 44]. In this paper, we study the impact of TTE on certification by experimenting with two well-known certified defenses: Randomized Smoothing [11] and SmoothAdv [44]. Our experiments find that TTE’s benefits also appear in certification, where TTE can boost these defenses by sizable margins.

3. Methodology

We study the impact of introducing TTE on adversarial robustness. We observe that traditional image transforms can provide simple transforms to be used in TTE. Thus, we aim at studying the impact that introducing customary im-

¹Code is available at <https://github.com/juancprzs/TTE>.

Algorithm 1 TTE Wrapper pseudo-code in PyTorch style.

```

# C: number of channels
# H: image height
# W: image width
# N: number of classes
class TTEWrapper:
    def init(self, model, transforms):
        # transforms: list of differentiable functions
        self.model = model
        self.transforms = [lambda x: x] # the identity
        self.transforms += transforms

    def forward(self, x):
        # apply transforms
        x = cat([t(x) for t in self.transforms])
        # move transforms to the batch dimension
        x = x.view(-1, C, H, W)
        # forward
        s = self.model(x)
        # move scores of transforms to other dimension
        s = s.view(len(self.transforms), -1, N)
        # average scores across transforms dimension
        return s.mean(dim=0)

```

cat: concatenation.

age transforms at test time has on adversarial robustness. In particular, we set to study a simple TTE wrapper for trained classification models. This wrapper is a module that receives an image, augments the image with a fixed set of transforms, inputs the image to the model, and, finally, ensembles the outputs through averaging.

Our selection of transforms for the wrapper is based on facilitating the assessment of adversarial robustness of the wrapped model. Specifically, we are interested in wrapping models and then properly conducting white-box attacks, since these attacks are at the core of the most challenging adversarial settings. The transforms we select, thus, must not hinder nor obstruct the computation of the gradient with respect to the model’s input. Hence, we select three well-known label-preserving transforms that are common in training routines: (i) horizontal flips, (ii) padding-and-cropping, and (iii) the composition of these two.

Note that these transforms are easily implemented in a **differentiable and deterministic** manner: all transforms are implemented by indexing the input tensor, which, during the backward pass, translates to simply directing the gradient to the selected indices. Further, we remark that we refrain from introducing stochasticity in the transforms, contrary to common implementations of training and testing routines. That is, the set of transforms and their parameters are kept *fixed* after initialization: flips are performed deterministically, and the sections being cropped are permanent. Once the set of transforms and their parameters have been set, we instantiate the transforms and use them for initializing the wrapper. Hence, we emphasize that our approach does not hinder nor obstruct computing the gradient of the wrapper’s output with respect to its input. Please refer to Algorithm 1 for pseudo-code of the wrapper, and to the **Supplementary Material** for the pseudo-codes of the transforms we study.

Table 1. **TTE and undefended models.** We test the impact of TTE on undefended models. TTE increases the adversarial robustness of undefended models by sizable margins. Best results in **bold**.

CIFAR10			
Method	Clean	Robust	Diff.
ResNet-18	92.58	16.18	+13.63
ResNet-18 + TTE	93.42	29.81	
CIFAR100			
ResNet-18	76.66	0.72	+0.96
ResNet-18 + TTE	77.57	1.68	

In the following section, we conduct a comprehensive empirical study of the adversarial robustness effects of equipping models with TTE. Remarkably, we find that TTE can consistently improve adversarial robustness across defenses, while requiring less than 30 lines of code.

4. Experiments

We conduct a comprehensive empirical study regarding the effects of TTE on adversarial robustness. We vary defenses, datasets, and attacks, while also gathering insights into TTE’s inner workings. We find that TTE yields sizable and consistent boosts in robustness, both against strong attacks and for certification purposes.

4.1. Adversarial Robustness Assessment

A reliable assessment of adversarial robustness is fundamental to our study. Hence, we use the challenging AutoAttack benchmark [13] for estimating adversarial robustness. AutoAttack is a parameter-free ensemble of diverse attacks that has shown outstanding capabilities of identifying vulnerabilities in adversarial robustness defenses. AutoAttack is composed of: AutoPGD, aiming at optimizing either Cross Entropy (CE) or a targeted version of a Difference-of-Logits-Ratio [13] (denoted as *APGD-CE* and *APGD-T*, respectively); targeted Fast Adaptive Boundary [12] (denoted as *FAB-T*) that aims at perturbation minimization and has shown promising results against gradient-masking; and Square Attack [1], a norm-bounded score-based black-box attack that does not rely on gradient information (denoted as *Square*). High performance against such a diverse and powerful ensemble has shown to provide an accurate assessment of adversarial robustness [13].

For all experiments, we run AutoAttack and report both clean accuracy and *Robust* accuracy, where the latter is defined as the per-instance worst case across all attacks. For the main experiments and with completeness in mind, we also report the accuracy against individual attacks composing AutoAttack.

4.2. TTE on Undefended Models

We begin our study by analyzing the effect of introducing TTE on the adversarial robustness of undefended mod-

els. We study this setting by conducting nominal training of a ResNet-18 [22] on CIFAR10 and CIFAR100 [27], and then attacking these models both with and without TTE. The attack strength values (ϵ) that the robustness community usually employs are capable of dropping the accuracy of undefended models to approximately 0% [7]. Thus, for this experiment, we use a weaker attack strength of $\epsilon = 2/255$ that allows us to observe variations in adversarial robustness. Table 1 summarizes the results. For both datasets, we find consistent increments in both clean and robust accuracy. In particular, robust accuracy increases by a remarkable 13% on CIFAR10, and by $\approx 1\%$ on CIFAR100. Our results thus provide evidence that TTE is capable of boosting the adversarial robustness of models without hurting accuracy on clean samples.

4.3. Boosting the State-of-the-Art

CIFAR. The AutoAttack benchmark hosts an online leaderboard², where current state-of-the-art (SOTA) defenses on various datasets are ranked according to their performance against AutoAttack. To evaluate the impact of equipping SOTA defenses with TTE, we download pre-trained models of defense approaches from the CIFAR10 leaderboard, add our TTE wrapper, and run AutoAttack on them with a perturbation budget of $\epsilon = 8/255$. Our selection of defenses for evaluation is based on high-performance in the AutoAttack benchmark and availability of the trained models (either from the official repository or upon contacting the respective authors). In particular, we test TTE on six high-performing defenses: Adversarial Training with Early Stopping (ATES) [48], Tradeoff-inspired Adversarial Defense via Surrogate-loss minimization (TRADES) [61], ImageNet pre-training (IN-Pret) [23], robust-network pruning (HYDRA) [45], Misclassification Aware Adversarial Training (MART) [53], Adversarial Weight Perturbation (AWP) [55], and the method of Gowal *et al.* [20]. Note that only three of these defenses are available for CIFAR100, as this dataset is much less studied for adversarial robustness.

We report the performance of the standard and TTE-enhanced versions of these defenses in Table 2. In line with previous works [28, 47, 22], our results show how equipping models with TTE generally enhances their clean accuracy. More importantly, we observe that enhancing defenses with TTE consistently increases the robustness of all defenses across both attacks and datasets. In CIFAR10, for instance, TTE increases robust accuracy by 2.04% on average. Analogously, in CIFAR100, we observe an average increase of 1.29%.

Note that the strongest defense for CIFAR10 is Gowal *et al.*, and for CIFAR100 the strongest defense is AWP. Equipping these top-performing defenses with TTE increases

their robust accuracy by 1.26% and 0.86% on CIFAR10 and CIFAR100, respectively.

ImageNet. Adversarial robustness defenses in ImageNet [15] are much less common than in CIFAR, mainly due to the computational costs associated with conducting adversarial attacks on such large images and vast amount of instances. A recent robustness assessment by Dong *et al.* [16] suggests that the SOTA defense on ImageNet is Feature Denoising (FD) [56], which combines large-scale adversarial training [32] with the introduction of denoising blocks into the model’s architecture. We download the pre-trained models from the official implementation and run AutoAttack on the standard and TTE-enhanced versions³. Following the methodology of [56], we run attacks with a perturbation budget of $\epsilon = 16/255$.

We report the results of this experiment in Table 3. These results show that, even in the large-scale regime of ImageNet, TTE is still able to provide improvements in adversarial robustness. In fact, we observe consistent increase in accuracy against every attack of the AutoAttack benchmark. In particular, we find that TTE is able to provide a remarkable improvement of 2.21% in robust accuracy. It is worthwhile to mention that we report the results corresponding to the best TTE selection, while leaving the extensive ablation results to the **supplementary material**.

4.4. Transform Selection

In Section 3, we established that the transforms considered in this study are horizontal flips, pad-and-crop, and the composition of these two. These transforms were selected so as to facilitate the assessment of adversarial robustness by introducing transforms that are both differentiable and deterministic. Since the padding-and-cropping operation is parameterized by the indices of the crop (the padding size is another parameter, but in our study we kept it fixed to 4), there is a large number of transforms that could be generated from an image. The combination of such transforms with horizontal flipping yields a vast search space. In this section, we study a simple set of these transforms, and show that a large majority of transforms already provides sizable gains in adversarial robustness.

As subjects of our study, we choose TRADES [61] for experiments on CIFAR10 and FD [56] for experiments on ImageNet. We use the wrapper introduced in Section 3 around both models (as obtained from their official implementations), and test various selections of transforms for initializing the wrapper. In particular, we consider: (i) flip, (ii) crops, (iii) flip + crops, and (iv) flip + crops + flipped-crops. Whenever we use a crop, we vary the number of extracted crops from one to four. For all the cases that in-

²Available at <https://github.com/fra31/auto-attack>

³Since this model was not trained on zero-padded images, we extract 224×224 crops from the 256×256 center crop of the resized image. We resize the shortest side to 256 while preserving the height-width ratio.

Table 2. **Adversarial robustness on CIFAR.** We compare the standard and the best TTE-enhanced versions of several defenses from the AutoAttack leaderboard, whose models are publicly available, on the CIFAR10 and CIFAR100 datasets. We report the clean accuracy, the adversarial accuracy against the individual attacks that compose AutoAttack, the worst-case (*Robust*) accuracy in (%), and the average difference across datasets. Best results in clean and robust accuracies are shown in **boldface**. Note that *all* defenses gain adversarial robustness when equipped with TTE.

	Method	Clean	APGD-CE	APGD-T	FAB-T	Square	Robust	Difference	Average
CIFAR10	ATES [48]	86.84	53.5	51.5	51.91	59.77	51.46	+2.71	+2.04
	ATES + TTE	86.86	56.48	54.19	54.70	60.67	54.17		
	TRADES [61]	84.92	55.31	53.12	53.55	59.41	53.11	+2.38	
	TRADES + TTE	85.14	57.46	55.51	55.88	60.22	55.49		
	IN-Pret [23]	87.11	57.65	55.32	55.68	62.40	54.92	+1.51	
	IN-Pret + TTE	87.13	59.06	56.44	56.73	63.14	56.43		
	MART [53]	87.50	62.18	56.80	57.34	64.87	56.75	+2.19	
	MART + TTE	87.79	63.95	58.94	59.51	65.62	58.94		
	HYDRA [45]	88.98	60.13	57.66	58.42	65.01	57.64	+2.74	
	HYDRA + TTE	88.82	62.82	60.40	60.91	66.03	60.38		
	AWP [55]	88.25	63.81	60.53	60.98	66.18	60.53	+1.46	
	AWP + TTE	88.07	64.95	61.99	62.52	66.48	61.99		
	Gowal <i>et al.</i> [20]	89.48	66.16	63.26	63.74	69.10	63.29	+1.26	
	Gowal <i>et al.</i> + TTE	89.41	67.19	64.55	64.88	69.29	64.55		
CIFAR100	ATES [48]	62.82	26.78	24.98	25.23	31.27	24.96	+1.83	+1.29
	ATES + TTE	63.47	28.9	26.8	27.15	32.21	26.79		
	IN-Pret [23]	59.37	33.45	29.03	29.34	34.55	28.61	+1.19	
	IN-Pret + TTE	59.38	33.96	29.59	29.87	34.86	29.50		
	AWP [55]	60.38	33.56	29.16	29.48	34.66	29.15	+0.86	
	AWP + TTE	60.39	34.11	30.03	30.26	34.64	30.01		

Table 3. **Adversarial robustness on ImageNet.** We compare the standard and the best TTE-enhanced versions of Feature Denoising [56] on the ImageNet dataset. We report the clean accuracy, the adversarial accuracy against individual attacks comprising AutoAttack, and the worst-case (*Robust*) accuracy in (%). Following the experimental setup in [56], we run this experiment with $\epsilon = 16/255$.

Method	Clean	APGD-CE	APGD-T	FAB-T	Square	Robust	Difference
FD [56]	65.32	7.91	4.31	7.87	23.76	4.23	+2.21
FD + TTE	65.87	9.29	6.53	9.23	26.34	6.44	

clude crops, we randomly *initialize* the indices of the crop, and keep them fixed after, *i.e.* note that the transform is, again, *deterministic*.

We report the evaluation of these transforms for TRADES in Table 4 and for FD⁴ in Table 5. Our results show that *any* of the studied transforms provides gains in adversarial robustness. Further, we note that even simple transforms already provide sizable gains. For instance, simply ensembling the original input with its flipped version boosts robust accuracies by 1.70% and 0.97% for TRADES and FD, respectively. For TRADES, we obtain the largest increase both in clean and robust accuracies with an ensemble composed of a flip, three crops, and three flipped-crops. This particular set of transforms achieves a remarkable boost of 2.38% in robust accuracy. However, we report the largest gain in robust accuracy (+1.98%) for FD when the image is processed jointly with four crops, four flipped-crops and the flipped instance. Our results show that TTE provides consistent gains in adversarial robustness across a

simple set of transforms for both benchmarks. This outcome suggests that, despite the large size of the space of possible transforms, finding a set of transforms that provides sizable gains in adversarial robustness is effortless. In particular, in our experiments, *all* the transforms provided gains in adversarial robustness.

4.5. Are all crops created equal?

In Section 4.4, we studied how several transforms affect adversarial robustness. From Tables 4 and 5, we notice an odd phenomenon: increasing the number of crops does not always yield larger robustness gains. Since the crop location is random, this phenomenon suggests that there are crops that provide larger robustness gains than others.

We conduct an exhaustive search over the possible crops, and record changes in clean and robust accuracy. As mentioned in Section 4.4, we fix the padding size to 4 in the pad-and-crop transform. Hence, the height and width of each image increases by 8 (each side increases by 4). Thus, the total number of possible crops is $(8 + 1)^2 = 81$.

We conduct this experiment on TRADES and evalu-

⁴Due to computational costs, we test FD with the official implementation’s 30-step PDG attack (PGD³⁰) under the ℓ_∞ norm with $\epsilon = 16/255$.

Table 4. **Adversarial robustness gains of various transforms on CIFAR10.** We test the impact in adversarial robustness of introducing various transforms to TTE on TRADES. We report clean and robust accuracies, and the difference in robustness between each TTE-enhanced model and TRADES. Robust accuracies larger than that of TRADES are shown in **boldface**.

Method	Clean	Robust	Diff.
TRADES	84.92	53.11	-
+ flip	85.07	54.81	+1.70
+ 1 crop	84.86	53.35	+0.24
+ 2 crops	84.85	53.35	+0.24
+ 3 crops	84.89	53.63	+0.52
+ 4 crops	84.87	53.34	+0.23
+ flip + 1 crop	85.09	55.08	+1.97
+ flip + 2 crops	84.82	54.68	+1.57
+ flip + 3 crops	85.07	54.76	+1.65
+ flip + 4 crops	85.11	54.43	+1.32
+ flip + 1 crop + 1 flipped-crop	85.13	55.19	+2.08
+ flip + 2 crops + 2 flipped-crops	85.06	55.24	+2.13
+ flip + 3 crops + 3 flipped-crops	85.14	55.49	+2.38
+ flip + 4 crops + 4 flipped-crops	85.13	55.49	+2.38

Table 5. **Adversarial robustness gains of various transforms on ImageNet.** We test the impact in adversarial robustness of introducing various transforms to TTE on Feature Denoising (FD) [56]. We report clean and PGD³⁰ accuracies, and the difference in robustness between each TTE-enhanced model and FD. PGD³⁰ accuracies larger than that of FD are shown in **boldface**.

Method	Clean	PGD ³⁰	Diff.
FD [56]	65.32	50.20	-
+ flip	65.38	51.17	+0.97
+ 1 crop	65.50	51.07	+0.87
+ 2 crops	65.51	50.84	+0.64
+ 3 crops	65.78	51.20	+1.00
+ 4 crops	65.74	51.21	+1.01
+ flip + 1 crop	65.56	51.69	+1.49
+ flip + 2 crops	65.59	51.77	+1.57
+ flip + 3 crops	65.81	51.80	+1.60
+ flip + 4 crops	65.76	51.43	+1.23
+ flip + 1 crop + 1 flipped-crop	65.69	51.47	+1.27
+ flip + 2 crops + 2 flipped-crops	65.68	51.36	+1.15
+ flip + 3 crops + 3 flipped-crops	65.87	51.88	+1.68
+ flip + 4 crops + 4 flipped-crops	65.85	52.17	+1.98

ate on CIFAR10. Since the evaluation of each model is computationally expensive, we refrain from using the official TRADES model (a large WideResNet-34-10 [58] model), and rather train a smaller ResNet-18 [22] model with the official TRADES implementation. This trained model achieves a clean accuracy of 80.96% and a robust accuracy of 48.64% against AutoAttack. For studying the impact of specific crops on robustness, we are interested in the evaluation of the model under two settings: (i) when the model’s input is *only* the selected crop, *i.e.* the model is not given the original image, and (ii) when the model’s input is both the original image *and* the crop, that is, TTE itself.

Table 6. **Accuracy under APGD-T attacks vs. optimization iterations and attack strength.** We study how accuracy under APGD-T attacks changes as we vary (i) iterations for optimization (top table), and (ii) attack strength (bottom table). Our results suggest that TTE does *not* induce gradient obfuscation.

Optimization iterations				
Iterations	5	10	50	100
TRADES	49.92	49.12	48.71	48.69
TRADES + TTE	52.11	51.54	51.41	51.40
Attack strength (ϵ)				
ϵ	8/255	16/255	32/255	64/255
TRADES	48.69	15.84	0.72	0.00
TRADES + TTE	51.40	18.85	0.95	0.01

Figure 5 displays our results as heatmaps. The x and y coordinates of the heatmap range from 0 through 8, and are interpreted as the offset from the top-left corner of the zero-padded version of the image. These heatmaps show appealing spatial patterns. We notice that when using only the cropped versions of the images, spatial translations closer to the origin provide the best performance in both clean and robust accuracy. In particular, cropped images can even achieve better clean and robust accuracy than their uncropped versions under most one-pixel shifts. However, when using the original image *and* its cropped version, the optimal crops shift towards the edges. In general, crops that arise from translations in both directions tend to consistently boost performance. Although a clear pattern in clean accuracy is elusive, robust accuracy tends to symmetrically improve as the selected crop moves towards the corners of the zero-padded image, with the exception of the corners themselves. This result implies that the model benefits from seeing the most shifted views of the image up to a certain threshold, but further shifts may be counterproductive.

4.6. Is TTE obfuscating the gradient?

In this paper, we construct the set of transforms to be studied based on a single criterion: that such transforms would ease the assessment of adversarial robustness. We consider transforms that would not hinder nor obstruct the computation of the gradient with respect to the input image. Thus, we study transforms that are differentiable and deterministic. Based on this criterion, we expect that adding the TTE wrapper does not lead to gradient obfuscation, which could yield an inaccurate assessment of adversarial robustness and, hence, a false sense of security [2]. Here, we empirically test this hypothesis. In particular, we are interested in knowing whether TTE prevents the model from producing “useful gradients” for iterative optimization attacks [38, 2]. Following [2], we check how the behavior of the model’s performance changes when the iterative attack’s parameters vary in terms of: (i) number of optimization steps and (ii) attack strength.

This experiment requires a large number of runs. Hence,

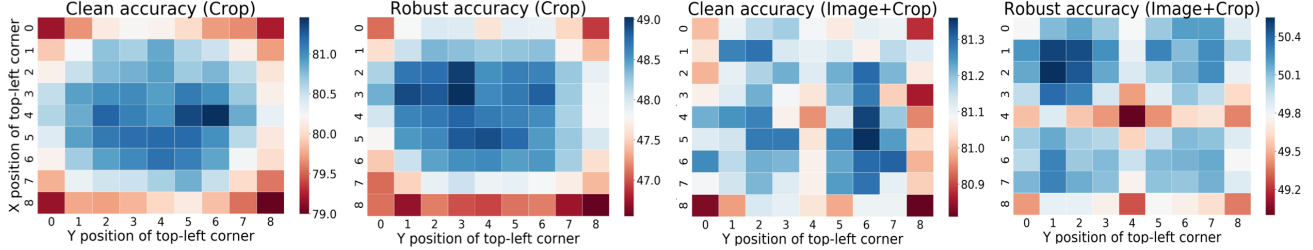


Figure 2. **Impact of specific crops on performance.** The heatmaps depict the clean and robust accuracies as a function of the location of the crops. The first two plots are obtained by only feeding cropped versions of CIFAR10 images, and the last two by feeding the original image *and* its cropped version. We assessed a total of 81 crops corresponding to all the possible translations in each scenario.

we follow the same experimental setup as in Section 4.5 and use a ResNet-18 model trained with TRADES. We compare this defense against its TTE-enhanced version that could cause, in theory, the largest effect on gradient obfuscation: original image + flip + 4 crops + 4 flipped-crops. Hence, we expect this experiment to provide an upper-bound to the (possible) gradient obfuscation effect of introducing TTE. We compute the adversarial accuracy of these defenses against the *APGD-T* [13] attack from the AutoAttack ensemble, which is a strong iterative gradient-based adversary. We vary the number of optimization steps from 1 to 100 (the default number of steps for *APGD-T* in AutoAttack), and the attack strength from $1/255$ to $64/255$.

We report our main results in Table 6, and detail the rest in the **supplementary material**. We find that TTE suffers both when increasing the attack’s number of iterations and its strength. Notably, the performance eventually reaches approximately 0% when $\epsilon = 64/255$. Combined with the observation that white-box attacks are more successful than black-box attacks in fooling an TTE-enhanced model (as shown in Tables 2 and 3), these results suggest that introducing TTE does *not* induce gradient obfuscation.

4.7. Matching train- and test-time transforms

Our experiments have considered three transforms: flip, pad-and-crop, and the combination of these two. We have shown how applying these transforms at *test* time boosts the robustness of several defenses. We note that *all* the defenses we studied also use these transforms at *train* time, as is common practice in the training routines for image classifiers. Thus, in all our experiments, a match has existed between transforms on which defenses were trained and the ones we introduced at test time. Here, we experiment with introducing a mismatch between the train- and test-time transforms. Specifically, we study two setups: (i) introducing a transform that is not seen during training (Gaussian filtering), and (ii) training the model by *removing* the pad-and-crop transform and then testing on padded-and-cropped images.

Gaussian filtering. We study the impact of introducing a Gaussian-filtering transform into TTE. We choose this transform based on spectral properties of adversarial exam-

Table 7. **Matching train- and test-time transforms.** We introduce a mismatch between train- and test-time transforms and record variations in performance. We induce a mismatch by (i) testing with TTE including Gaussian-filtering (top sub-table), and (ii) training a TRADES model without the pad-and-crop transformation (TRADES^{nc}) and testing with TTE including pad-and-crop transforms (bottom sub-table). Results on CIFAR10 show that TTE’s boosts require matching train- and test-time transforms.

Method	Clean	Robust	Diff.
TRADES	84.92	53.11	-
+ Gaussian ($k = 3, \sigma = 1$)	81.59	50.21	-2.90
+ Gaussian ($k = 3, \sigma = 2$)	81.27	49.63	-3.48
+ Gaussian ($k = 5, \sigma = 1$)	80.36	49.19	-3.92
+ Gaussian ($k = 5, \sigma = 2$)	76.78	45.71	-7.40
TRADES ^{nc}	82.40	48.64	-4.47
TRADES ^{nc} + 1 crop	81.28	47.04	-6.07
TRADES ^{nc} + 2 crops	77.03	42.23	-10.88
TRADES ^{nc} + 3 crops	80.66	46.57	-6.54
TRADES ^{nc} + 4 crops	77.79	41.39	-11.72

ples, as recent works find a relation between high-frequency components and adversarial examples [52, 57]. These findings suggest that using Gaussian filtering, a low-pass filter, could increase the adversarial robustness of a defense. Note that this transform can also be implemented in a differentiable manner, complying with our requirements.

Thus, we modify TTE to only include Gaussian filtering, equip TRADES with TTE, and conduct attacks. We consider Gaussian filters with filter sizes $k \in \{3, 5\}$, and standard deviations $\sigma \in \{1, 2\}$. We report the results in Table 7 (top sub-table). Results indicate that introducing Gaussian filtering at test time, when the defense was *not* trained on such transforms, is *detrimental* to adversarial robustness.

Removing pad-and-crop from training. We train a TRADES model by removing the usual pad-and-crop transform, and refer to this model as TRADES^{nc}. We record the performance of TRADES^{nc} without TTE and with TTE based only on pad-and-crop transforms. We vary the number of crops from one to four. We report results in Table 7 (bottom sub-table). We note that (i) the pad-and-crop train-time transform is important for TRADES: removing this transform decreases clean and robust accuracies, approximately, by 2%; and (ii) adding a padded-and-cropped ver-

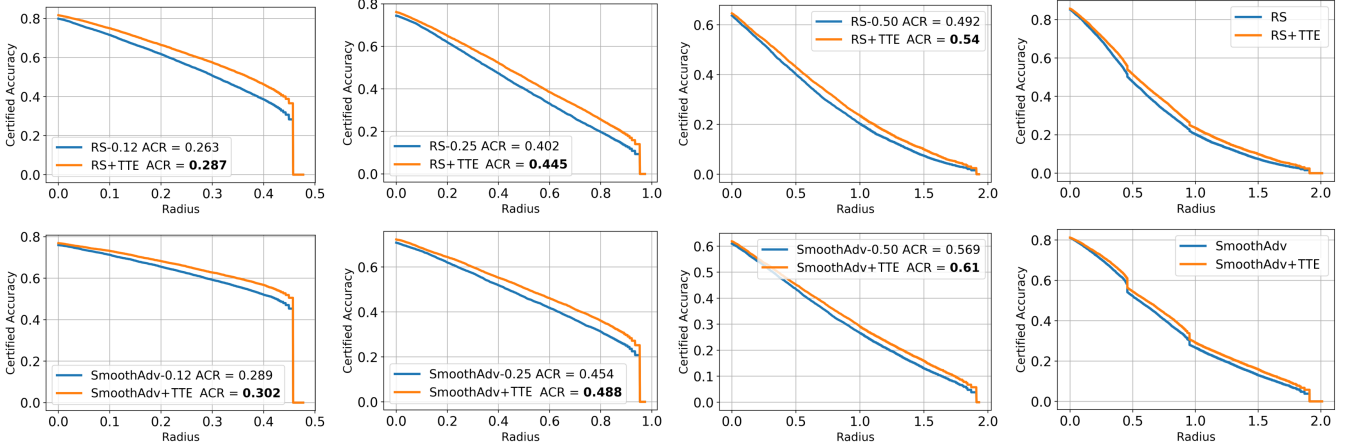


Figure 3. **Boosting certified defenses with TTE.** We introduce TTE into two certified defenses: Randomized Smoothing (RS) [11] and SmoothAdv [44] (top and bottom row, respectively). Each plot reports a certified accuracy curve. The first three columns show certified accuracy for $\sigma \in \{0.12, 0.25, 0.5\}$, respectively, and each legend reports the Average Certified Radius (ACR). The fourth column shows the envelope of the corresponding first three plots: the largest certificate that can be granted to each method across σ values.

sion of the image is detrimental to the adversarial robustness of TRADES^{nc}: any number of additional transformed versions of the image hurt both clean and robust accuracies. Our results suggest that, for a TRADES model, there is large distribution mismatch between the original images and their padded-and-cropped versions.

The results in Table 7 indicate that matching train- and test-time transforms is fundamental for adversarial robustness. That is, simulating TTE through image transforms is beneficial for adversarial robustness, if and only if, a match exists between train- and test-time transforms.

4.8. Boosting Certified Defenses

Our experiments show that TTE provides improvements in the empirical assessment of adversarial robustness. Here, we test whether these improvements also transfer to the domain of certifiably robust models. Formally, for an input x with true label y , a model f is certifiably robust at x with radius R if $f(x) = f(x + \delta) = y \quad \forall \|\delta\|_p \leq R$. Thus, the certified accuracy of a model f at a radius R is defined as the portion of the test set for which the model is certifiably robust with a radius of at least R . Recently, two closely-related methods have been proposed for certification. Cohen *et al.* [11] proposed Randomized Smoothing (RS), a technique that provides tight certified classification to an input by assigning the most probable class predicted by f when x is exposed to Gaussian noise of the form $\mathcal{N}(0, \sigma^2 I)$. Moreover, Salman *et al.* [44] proposed SmoothAdv, extending RS with adversarial training to boost certification.

We study TTE’s effect on certification on CIFAR10. We follow [11] for RS and train a Gaussian-augmented ResNet-18. For SmoothAdv, we follow [44] and apply Gaussian augmentation on PGD-generated images. We study certification in the ℓ_2 sense and test with $\sigma \in \{0.12, 0.25, 0.5\}$.

We certify TTE-enhanced models on the test set ($\mathcal{S}_{\text{test}}$) and report both certification curves and the Average Certified Radius [60], $\text{ACR} = 1/|\mathcal{S}_{\text{test}}| \sum_{(x,y) \in \mathcal{S}_{\text{test}}} R(x) \mathbb{1}_{\{f(x)=y\}}$, where $\mathbb{1}$ is the indicator function.

We summarize our results in Figure 3. We report a curve per σ and the curves’ envelope, *i.e.* the largest certificate that can be granted to each defense across all σ values. We observe that TTE consistently improves certified accuracy for all radii, for all values of σ , and for both defenses. In particular, at $\sigma = 0.5$, TTE improves the ACR by approximately 10% and 7% on RS and SmoothAdv, respectively. We remark that equipping these certified defenses with TTE comes at no cost during training, as TTE is only used at test time. That is, similar to our experiments on the empirical assessment of adversarial robustness, TTE provides boosts in the certification domain without requiring re-training.

Our results demonstrate that the benefits of TTE are not confined to the empirical assessment of robustness: TTE can also enhance *certified* adversarial robustness.

5. Conclusions

In this work, we analyzed the effect of using Test-time Transformation Ensembling (TTE) on adversarial robustness. We conducted a comprehensive empirical study on the adversarial robustness effects of leveraging TTE through customary image transforms. Our results demonstrate that TTE is a simple yet effective technique for improving adversarial robustness. Notably, we showed that the performance of several SOTA adversarial robustness defenses can be boosted by including TTE, and, further, that these benefits transfer to the domain of certified defenses.

Acknowledgments. This work was partially supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research.

References

- [1] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: a query-efficient black-box adversarial attack via random search. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2, 3
- [2] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International Conference on Machine Learning (ICML)*, 2018. 2, 6
- [3] Anish Athalye, Logan Engstrom, Andrew Ilyas, and Kevin Kwok. Synthesizing robust adversarial examples. In *International conference on machine learning (ICML)*, 2018. 2
- [4] Sven Bambach, David Crandall, Linda Smith, and Chen Yu. Toddler-inspired visual object learning. In *Advances in neural information processing systems (NeurIPS)*, 2018. 1
- [5] J.G. Bremner. *Infancy*. Wiley, 1994. 1
- [6] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, Aleksander Madry, and Alexey Kurakin. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019. 2
- [7] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, 2017. 1, 4
- [8] Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C Duchi, and Percy S Liang. Unlabeled data improves adversarial robustness. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 2
- [9] Jinghui Chen, Yu Cheng, Zhe Gan, Quanquan Gu, and Jingjing Liu. Efficient robust training via backward smoothing. *arXiv preprint arXiv:2010.01278*, 2020. 2
- [10] Tianlong Chen, Sijia Liu, Shiyu Chang, Yu Cheng, Lisa Amini, and Zhangyang Wang. Adversarial robustness: From self-supervised pre-training to fine-tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [11] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning (ICML)*, 2019. 2, 8
- [12] Francesco Croce and Matthias Hein. Minimally distorted adversarial examples with a fast adaptive boundary attack. In *International conference on machine learning (ICML)*, 2020. 3
- [13] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International Conference on Machine Learning (ICML)*, 2020. 2, 3, 7
- [14] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2019. 2
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 2, 4
- [16] Yinpeng Dong, Qi-An Fu, Xiao Yang, Tianyu Pang, Hang Su, Zihao Xiao, and Jun Zhu. Benchmarking adversarial robustness on image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 4
- [17] Yinpeng Dong, Hang Su, Baoyuan Wu, Zhifeng Li, Wei Liu, Tong Zhang, and Jun Zhu. Efficient decision-based black-box adversarial attacks on face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [18] Justin Gilmer, Nicolas Ford, Nicholas Carlini, and Ekin Cubuk. Adversarial examples are a natural consequence of test error in noise. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019. 1
- [19] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2015. 1, 2
- [20] Sven Gowal, Chongli Qin, Jonathan Uesato, Timothy Mann, and Pushmeet Kohli. Uncovering the limits of adversarial training against norm-bounded adversarial examples, 2020. 4, 5
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015. 1
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2016. 2, 4, 6
- [23] Dan Hendrycks, Kimin Lee, and Mantas Mazeika. Using pre-training can improve model robustness and uncertainty. In *International Conference on Machine Learning (ICML)*, 2019. 2, 4, 5
- [24] Swapna Jayaraman and Linda B Smith. Faces in early visual environments are persistent not just frequent. *Vision research*, 2019. 1
- [25] Ildoo Kim, Younghoon Kim, and Kim Sungwoong. Learning loss for test-time augmentation. In *Advances in neural information processing systems (NeurIPS)*, 2020. 2
- [26] Minseon Kim, Jihoon Tack, and Sung Ju Hwang. Adversarial self-supervised contrastive learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [27] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. In *University of Toronto, Canada*, 2009. 4
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems (NeurIPS)*, 2012. 1, 2, 4
- [29] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *International Conference on Learning Representations Workshop (ICLRW)*, 2017. 2
- [30] Yi Li, Haozhi Qi, Jifeng Dai, Xiangyang Ji, and Yichen Wei. Fully convolutional instance-aware semantic segmentation.

- In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [31] Xuanqing Liu, Minhao Cheng, Huan Zhang, and Cho-Jui Hsieh. Towards robust neural networks via random self-ensemble. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [32] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations (ICLR)*, 2018. 1, 2, 4
- [33] Maureen E McQuillan, Linda B Smith, Chen Yu, and John E Bates. Parents influence the visual learning environment through children’s manual actions. *Child development*, 2020. 1
- [34] Tomas Mikolov, Kai Chen, Greg S. Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space, 2013. 1
- [35] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. 2013. 1
- [36] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2016. 2
- [37] Nikita Moshkov, Botond Mathe, Attila Kertesz-Farkas, Reka Hollandi, and Peter Horvath. Test-time augmentation for deep learning-based cell segmentation on microscopy images, 2020. 2
- [38] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, 2017. 6
- [39] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 2
- [40] Lynn K Perry, Larissa K Samuelson, Lisa M Malloy, and Ryan N Schiffer. Learn locally, think globally: Exemplar variability supports higher-order generalization and word learning. *Psychological science*, 2010. 1
- [41] J. Piaget and B. Inhelder. *The Psychology Of The Child*. Basic Books, 2008. 1
- [42] Edward Raff, Jared Sylvester, Steven Forsyth, and Mark McLean. Barrage of random transforms for adversarially robust defense. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [43] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. Certified defenses against adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [44] Hadi Salman, Jerry Li, Ilya Razenshteyn, Pengchuan Zhang, Huan Zhang, Sebastien Bubeck, and Greg Yang. Provably robust deep learning via adversarially trained smoothed classifiers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 2, 8
- [45] Vikash Sehwal, Shiqi Wang, Prateek Mittal, and Suman Jana. Hydra: Pruning adversarially robust neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2, 4, 5
- [46] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Rob Fergus, and Yann LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2014. 2
- [47] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015. 2, 4
- [48] Chawin Sitawarin, Supriyo Chakraborty, and David Wagner. Improving adversarial robustness through progressive hardening. *arXiv preprint arXiv:2003.09347*, 2020. 4, 5
- [49] Lauren K Slone, Linda B Smith, and Chen Yu. Self-generated variability in object images predicts vocabulary growth. *Developmental science*, 2019. 1, 2
- [50] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [51] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*, 2014. 1
- [52] Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P Xing. High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 7
- [53] Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. Improving adversarial robustness requires revisiting misclassified examples. In *International Conference on Learning Representations (ICLR)*, 2019. 2, 4, 5
- [54] Eric Wong and Zico Kolter. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning (ICML)*, 2018. 2
- [55] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2, 4, 5
- [56] Cihang Xie, Yuxin Wu, Laurens van der Maaten, Alan L. Yuille, and Kaiming He. Feature denoising for improving adversarial robustness. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 4, 5, 6
- [57] 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 (NeurIPS)*, 2019. 7
- [58] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *British Machine Vision Conference (BMVC)*, 2016. 6
- [59] S. Zagoruyko, A. Lerer, T.-Y. Lin, P. O. Pinheiro, S. Gross, S. Chintala, and P. Dollár. A multipath network for object detection. In *British Machine Vision Conference (BMVC)*, 2016. 2
- [60] Runtian Zhai, Chen Dan, Di He, Huan Zhang, Boqing Gong, Pradeep Ravikumar, Cho-Jui Hsieh, and Liwei Wang. Macer: Attack-free and scalable robust training via maximizing certified radius. *International Conference on Learning Representations (ICLR)*, 2020. 2, 8
- [61] 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 *International Conference on Machine Learning (ICML)*, 2019. 2, 4, 5