This ICCV workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# **On the Effect of Pruning on Adversarial Robustness**

Artur Jordao and Helio Pedrini Institute of Computing, University of Campinas (UNICAMP), Brazil

arturjlcorreia@liv.ic.unicamp.br, helio@ic.unicamp.br

## Abstract

Pruning is a well-known mechanism for reducing the computational cost of deep convolutional networks. However, studies have shown the potential of pruning as a form of regularization, which reduces overfitting and improves generalization. We demonstrate that this family of strategies provides additional benefits beyond computational performance and generalization. Our analyses reveal that pruning structures (filters and/or layers) from convolutional networks increase not only generalization but also robustness to adversarial images (natural images with content modified). Such achievements are possible since pruning reduces network capacity and provides regularization, which have been proven effective tools against adversarial images. In contrast to promising defense mechanisms that require training with adversarial images and careful regularization, we show that pruning obtains competitive results considering only natural images (e.g., the standard and low-cost training). We confirm these findings on several adversarial attacks and architectures; thus suggesting the potential of pruning as a novel defense mechanism against adversarial images.

#### 1. Introduction

Despite achieving remarkable results in image classification [42, 25], modern and high-performance convolutional networks are easily fooled by adversarial images [16, 19]. Unlike natural (clean) images, adversarial images have their content modified to force a network to wrong its prediction. Figure 1 (left) shows an adversarial image and its effect on an over-parameterized network (leftmost). Since adversarial images exist in real-world settings, increasing the adversarial robustness of convolutional networks plays a role in safety- and security-critical applications.

Many mechanisms provide defense against adversarial images, for example, transfer learning [17, 39], regularization [51, 49] and data augmentation [6, 18, 27]. Regarding the latter, there is a successful category that combines clean and adversarial images during the training phase,



Figure 1. Predictive behavior of an over-parameterized network and its pruned counterpart on adversarial images. In this example, adversarial images are crafted by Fast Gradient Sign Method (FGSM). Left. Forwarding an adversarial image through the highcapacity (i.e., over-parameterized) network results in a wrong prediction - the network classifies the bird image as a panda. After removing its neurons, the network classifies the image correctly, suggesting that the pruned network became less vulnerable to the attack. Right. Accuracy on adversarial images (robustness) of ResNet56 and its pruned versions. On the standard FGSM attack, in which a network is attacked by adversarial images crafted from its parameters, the pruned ResNet56 (blue curve) achieves better robustness than its unpruned (red curve) counterpart. The robustness is even better when the adversarial images are yielded by ResNet56 and used to attack its pruned version, i.e., the attack is unaware of pruning (black curve). Note that such findings are confirmed under different attack severities.

named adversarial training [33]. Despite providing remarkable results, adversarial training requires careful regularization; otherwise, it hurts predictive ability on clean images [6, 36, 37, 49, 51]. Moreover, adversarial training might lack effectiveness when training data is scarce [36, 4] or the settings of the attack are changed [46].

Regardless of the defense mechanisms, most efforts attempt to circumvent the dilemma between adversarial robustness (accuracy on adversarial images) and generalization (accuracy on clean images), which means improving network accuracy on both domains.

According to previous works [43, 24, 20, 58], the success of adversarial images is due to the excessive overparameterization and capacity of deep networks. For example, Kaya et al. [24] and Hu et al. [20] observed that stopping inference in early layers improves adversarial robustness. Analogously, this mechanism aims at finding a shallow (low capacity) network inside a deeper network (high capacity). Another evidence is that networks with strong regularization are less vulnerable to adversarial images [6]. Building upon these findings, a natural idea is to increase adversarial robustness by reducing network capacity and forcing regularization while preserving generalization. In this work, we show that it is possible to satisfy such conditions by eliminating structures (filters, layers or both) from convolutional networks.

Our analysis relies on the evidence of previous research on compression and acceleration of deep networks, which have demonstrated the potential of pruning as a form of regularization capable of reducing overfitting and improving generalization [21, 26, 7, 2]. We leverage such behavior through the lens of adversarial images and demonstrate that removing filters and/or layers increases adversarial robustness without hurting generalization. Therefore, this process promotes an effective defense mechanism.

Contributions. Our key contribution is to demonstrate that eliminating structures of convolutional networks - pruning - increases their adversarial robustness, see Figure 1 (right). In order to confirm this, we prune different structures composing a convolutional network such as filters, layers and both. Importantly, throughout the pruning process, we take into account only clean images. In practice, this means that we can ignore the settings or any additional assumption of the attack when designing a defense mechanism. Further, it implies that we can avoid careful regularization and use the *low-cost* (standard – clean images only) training as a defense mechanism. By pruning layers, we show that shallower networks might obtain superior adversarial robustness than their deeper counterpart. This evidence exhibits an opposite phenomenon to a recent finding, which states that some defense mechanisms require deeper (costly) networks [51, 49]. All these findings appear without decreasing accuracy on clean images, which demonstrates the potential of pruning in satisfying the dilemma between generalization and robustness. However, assuming a negligible drop in generalization (to be precise less than one percentage point), we demonstrate that removing a single structure (e.g., only one filter) without any parameter adjustment enables a pruned network to obtain better robustness than its unpruned (defenseless) version.

We validate the claims above on several adversarial attacks (semantic-preserving, Fast Gradient Sign Method and occlusion), convolutional architectures (VGG, MobileNets and Residual networks) and datasets (CIFAR-10 and ImageNet). Additionally, we assess the effectiveness of pruning on many settings such as the structure removed, the pruning criterion and the scheme to adjust the parameters of the networks. On these configurations, pruning obtained gains in adversarial robustness of up to 4.29 percentage points. These gains substantially increase when we craft adversarial images using a convolutional architecture and attack a pruned version of other architecture. Compared to stateof-the-art defense mechanisms, which often involve training from scratch on adversarial images, pruning obtained competitive results performing few epochs of fine-tuning on clean images. Furthermore, it achieved one of the best trade-offs between robustness and generalization.

## 2. Preliminaries and Problem Statement

**Notations.** Define X and Y the set of images and their respective class labels. Let  $x \in X$  and  $y \in Y$  be a single image and label of X and Y, respectively. Let  $\mathcal{F}$  be a convolutional network that receives an input image x and outputs j probabilities, each one representing the probability of x belonging to the *jth* class. Thus, the classification of x using  $\mathcal{F}$  is  $\max(\mathcal{F}(x))$ , and  $\mathcal{F}$  correctly classifies x if  $\max(\mathcal{F}(x)) = y$ . For simplicity, we will omit the notation  $\max(.)$ .

Adversarial Attack. For a given image x, assume that a network  $\mathcal{F}$  correctly classifies x,  $\mathcal{F}(x) = y$ . An adversarial attack crafts an adversarial version of x, x', so that  $\mathcal{F}$  makes a misclassification, which means  $\mathcal{F}(x') \neq y$ . Equation 1 expresses the process of crafting an adversarial image.

$$x' = x + \alpha \delta, \delta \in \Delta,\tag{1}$$

where  $\delta$  represents a single corruption given a set of possible corruptions  $\Delta$ , which include semantic-preserving transformations and corruptions designed using the own parameters of a network (i.e., model-specific attacks). The parameter  $\alpha$  controls the severity/intensity of the corruption, for which low values become x' similar to x and visually imperceptible to humans, but the attack is less effective.

**Pruning Structures.** Let c be a criterion (a.k.a pruning criterion) that receives as input a network  $\mathcal{F}$  and assigns the importance for each structure (filters or layers) composing  $\mathcal{F}$ . Define such an operation as  $c(\mathcal{F})$ . Let S be a set of scores (i.e., importance) resulting from applying  $c(\mathcal{F})$ , with each element of S corresponding to the importance of a single structure of  $\mathcal{F}$ . Define I a set of indices indicating the p% lowest score (unimportant) structures of  $\mathcal{F}$  based on S; thereby,  $|I| = \lceil p \times S \rceil$ . The pruning algorithm removes the structures of  $\mathcal{F}$  is a thinner (removing filters) or shallower (removing layers) version of  $\mathcal{F}$ . After removing structures, the algorithm adjusts the parameters of the resulting network,  $\mathcal{F}'$ , for some epochs (fine-tuning).

One common practice is to remove the structures of  $\mathcal{F}$  iteratively. Such a strategy consists of using  $\mathcal{F}'$  of the previous iteration as input to the next iteration. This iterative nature implies that the pruned network of the *k*-th iteration has

Algorithm 1 Pruning Structures of Convolutional Networks
<b>Input:</b> Convolutional Network $\mathcal{F}$
<b>Input:</b> Number of Iterations K
Input: Pruning Criterion c
Input: Pruning Ratio p
<b>Output:</b> Pruned Convolutional Network $\mathcal{F}'$
for $k = 1$ to $K$ do
$S \leftarrow c(\mathcal{F}) \triangleright$ Assigns importance for each structure
$I \leftarrow p\%$ unimportant structures based on S
$\mathcal{F}' \leftarrow \mathcal{F} \setminus I \triangleright$ Removes the structures indexed by $I$
Update $\mathcal{F}'$ via standard fine-tuning (clean images)
$\mathcal{F} \leftarrow \mathcal{F}'$
end for

fewer structures than its preceding versions. Algorithm 1 summarizes the pruning process.

**Problem Statement.** Define  $Acc_{clean}(.)$  and  $Acc_{adv}(.)$  as the metrics that compute the accuracy of a convolutional network on clean (generalization) and adversarial (robustness) images, respectively. In this work, we verify the following statement:

 $Acc_{adv}(\mathcal{F}') > Acc_{adv}(\mathcal{F}), Acc_{clean}(\mathcal{F}') \approx Acc_{clean}(\mathcal{F}).$ 

The left part of the statement above allows us to answer the following research question: *Do pruned networks inherit the vulnerability to adversarial images of their unpruned counterpart?* Both sides allow us to answer the following: *Are pruned networks capable of improving robustness while preserving generalization?* 

### 3. Related Work

Adversarial Attacks. Adversarial attacks are mechanisms that add visually imperceptible perturbations to natural (clean) images focusing on degenerating the predictive ability of a network. These attacks may have access to the network parameters and architecture design [10, 33], or only to the network output [38]. For example, attacks such as Fast Gradient Signed Method (FGSM) and Projected Gradient Descent (PGD) use the gradient of loss to perturb clean images. Without access to network details (i.e., modelagnostic), it is also possible to make a network wrong its predictions. As demonstrated by Hendrycks et al. [16], semantic-preserving transformations (e.g., blur and contrast) applied to clean images are capable of degenerating the performance of high-capacity networks. According to previous works [55, 6, 40], even simple occlusions in clean images confound the network prediction.

Adversarial Defenses. Defense mechanisms are techniques that aim at reducing the vulnerabilities of convolutional networks to adversarial attacks. An effective defense mechanism is to increase the variability of training samples by data augmentation [55, 6, 18, 28]. From an empirical

perspective, Chun et al. [6] demonstrated that data augmentation techniques such as Cutout and Mixup are able to improve adversarial robustness. From a theoretical perspective, Zhang et al. [57] proved that Mixup minimizes the adversarial loss on adversarial images and operates as a form of regularization on clean images. Thus, it can theoretically improve both robustness and generalization.

The seminal work of Geirhos et al. [9] observed that convolutional networks trained on ImageNet exhibit a texture bias that decreases their representative performance. They demonstrated that augmenting training samples with style transfer (replacing the texture from an image with another) alleviates texture bias and increases adversarial robustness. Improving upon this idea, Shi et al. [40] mitigated the texture bias of convolutional networks by removing local textures from the layer's output (i.e., feature maps). For this purpose, the authors adopted a dropout-based approach in which flat and high-frequency regions of a feature map have a higher probability to be dropped (zeroed out). From extensive experiments, Mummadi et al. [34] disproved that removing texture bias increases adversarial robustness. Instead, they confirmed that robustness comes from the data augmentation by the style transfer. Li et al. [28] showed that debiasing only texture or shape information degrades network performance on clean images. To improve generalization and robustness, they proposed a label assignment strategy that explores both texture and shape.

Closely related to data augmentation, increasing data training by combining clean and adversarial images, named adversarial training, is a successful defense mechanism [33, 35]. It is worth mentioning that Cutout, MixUp and their variations naturally implement adversarial training on the occlusion attack, as they perturb regions of the image.

Despite the positive results, adversarial training requires careful regularization to avoid performance degradation on clean images [6, 37, 49, 51, 1]. Recent investigations revealed that adversarial training leads to inaccurate estimation of the batch normalization layers [51, 49, 46], which increases the gap between generation and robustness. As demonstrated by Xie et al. [49], this problem takes place due to mismatching distribution between clean and adversarial images. Furthermore, Xie and Yuille [51] demonstrated that adversarial training demands very deep networks to obtain better robustness. Since deeper models are computationally inefficient, obtaining adversarial robustness by these mechanisms might limit applicability on resource-constrained devices. Fortunately, we show an opposite phenomenon to Xie and Yuille [51], in which shallower networks (produced through removing layers) can provide gains in robustness.

In contrast to adversarial training and other data augmentation techniques, we show that it is possible to improve robustness considering only clean images. Thereby, we can avoid assumptions of the attack and careful regularization. Previously, Xie et al. [50] reached this advantage by adding random scale and padding to an image before forwarding it through the network. Unfortunately, their strategy incurs a higher computational cost since expanding scale is one factor that negatively impacts the computational performance of a network [42, 12].

**Pruning Structures.** Pruning consists of locating and removing the least important structures (filters or layers) from convolutional networks preserving their predictive ability as much as possible. To satisfy such conditions, previous works have demonstrated promising results using datadriven and data-agnostic criteria [31, 41, 32, 23, 29, 22, 5, 11]. While the first measures the importance of a structure according to its relationship with the data, e.g., using its features maps or loss, the latter estimates importance directly on the network parameters.

Most research on pruning aims at removing small structures such as filters or even weights. Veit et al. [45] observed that residual architectures exhibit negligible accuracy drop when removing some of their layers<sup>1</sup>. Building upon this, recent strategies are concentrated on removing entire layers [47, 48, 56, 44, 7]. Importantly, this family of pruning considers residual-based architectures only, as the evidence by Veit et al. [45] is invalid to plain networks, e.g., VGG.

Regardless of the criterion for assigning importance or structure removed, pruning has been confirmed as a form of regularization [21, 26, 7, 2], which reduces overfitting and improves generalization. In this work, we leverage this behavior through the lens of adversarial images, demonstrating its effect on robustness. Concurrently to our work, Ye et al. [53] demonstrated that adversarial training jointly with pruning leads to adversarial robustness. Hence, their work inherits all the drawbacks of adversarial training, for example, careful regularization and retraining from scratch when the attack is changed. On the other hand, our analysis prevents adversarial training, as it takes into account only clean images. Moreover, we conduct a more comprehensive evaluation since we investigate several pruning settings such as different schemes to parameters adjustment, structures removed, and pruning criteria.

## 4. Experiments

**Experimental Setup.** Unless stated otherwise, we prune the networks using one iteration of pruning (k = 1 in Algorithm 1). We shall see that multiple iterations of pruning marginally influence the relationship between robustness and generalization. To identify unimportant structures, we mainly focus on the Partial Least Squares (PLS) criterion [23]. However, we also demonstrate the influence of

pruning on robustness when the pruning algorithm considers other criteria. After pruning is complete, we fine-tune the resulting architecture for 200 and 12 epochs on CIFAR-10 and ImageNet, respectively. In this step, we use random crop and horizontal random flip as data augmentation [13]. We also employ this data augmentation during the training phase (i.e., before running the pruning). Such a setting plays an important role since it ensures we see adversarial images during the evaluation only.

To generate adversarial images, we employ the Fast Gradient Sign Method (FGSM) and model-agnostic attacks. In the latter, we use the semantic-preserving transformations (semantic for short) proposed by Hendrycks et al. [16] and the occlusion attack for which we fill the image center with a square matrix of zeros, as suggested by previous works [6, 55]. Unless otherwise specified, we consider the highest level of severity to the semantic-preserving ( $\delta = 4$ ) and occlusion ( $\delta = 16$ ), attacks and  $\delta = 8/255$  to FGSM. We explore these attacks on plain (VGG16), lightweight (MobileNetV2) and residual architectures (ResNet56 and ResNet50). We report the performance in robustness and generalization using the difference, in percentage points (p.p), between the unpruned and pruned network. Thus, positive values, (+), indicate gains while negative values, (-), indicate degradation. All experiments were run on a machine with 16GB RAM and a single NVIDIA GTX 1080.

Robustness and Generalization from Pruning. Our first experiment analyzes the effect of pruning on the trade-off between generalization (accuracy on clean images) and adversarial robustness (accuracy on adversarial images). Table 1 summarizes the results on CIFAR-10. From this table, we observe that the networks become more robust to adversarial images after removing some of their structures (filters, layers or both). For example, after pruning the filters of ResNet56, its robustness improved up to 3.68 p.p. In particular, considering all the attacks we evaluate, pruning improved at least 1.58 p.p the robustness of ResNet56. The same trend occurred when pruning its layers, but with a smaller improvement in adversarial robustness. More concretely, pruning this structure, the gains in robustness ranged from 1.05 p.p and 3.20 p.p. It is worth noting that pruning provided these gains without degrading network generalization. On the other hand, a different behavior takes place when we prune both filters and layers. On semanticpreserving, pruning both structures decreased robustness by 4.13 p.p. We believe that the regulation severity caused by pruning both structures of ResNet56 impaired its performance. In summary, despite the positive results on other attacks, the average improvement in pruning both structures is no better than removing filters and layers separated.

We observe that pruning MobileNetV2 and VGG16 is also an effective way of increasing their robustness. By comparing the gains of pruning these architectures to

<sup>&</sup>lt;sup>1</sup>For simplicity, we use the term layers, but the work by Veit et al. [45] confirmed such behavior when removing residual blocks – sets of layers where the input of the first layer is added to the output of the last layer [13].

Table 1. Difference between accuracy of the unpruned (defenseless) network and its pruned counterpart. The symbols (+) and (-) indicate improvement and degradation, respectively, regarding the unpruned network. Last column shows the average improvement. Overall, the process of removing structures from convolutional networks provides robustness without degrading generalization, which reveals that pruning is capable of satisfying the dilemma between these metrics.

Architecture	Structure	Semantic	Occlusion	FGSM	Clean	Average
	Filters	(+) 1.58	(+) 2.76	(+) 3.68	(+) 0.60	(+) 2.15
ResNet56	Layers	(+) 1.05	(+) 1.06	(+) 3.20	(+) 0.84	(+) 1.53
	Both	(-) 4.13	(+) 4.62	(+) 0.36	(-) 0.60	(+) 0.06
	Filters	(-) 0.60	(+) 3.35	(+) 0.64	(+) <b>0.37</b>	(+) 0.94
MobileNetV2	Layers	(-) 0.49	(+) 2.12	(+) 1.44	(+) 0.15	(+) 0.80
	Both	(+) <b>0.07</b>	(+) 2.56	(+) 1.05	(+) 0.17	(+) <b>0.96</b>
VGG16	Filters	(-) 1.0	(+) 4.86	(-) 2.21	(+) 0.89	(+) 0.63

ResNet56, we note that the average improvement decreased as a function of network depth. For example, the highest average improvement achieved by pruning ResNet56 (56-layer deep) was 2.15 p.p., while on MobileNetV2 (51layer depth) and VGG16 (16-layer deep) this improvement was 0.96 and 0.63 p.p. As suggested by previous works [47, 15, 14], pruning brings better benefits to overparameterized (i.e., deeper) networks. Our results corroborate this finding in the context of adversarial robustness.

Overall, the empirical evidence above shows that pruning improves robustness without sacrificing generalization, which is a desirable condition to any defense mechanism. Importantly, the pruning procedure takes into account clean images only. In contrast to previous research [51, 49], which states that effective defense mechanisms require deeper networks, our results by pruning layers suggest an opposite phenomenon – shallow networks can be as robust as their (unpruned) deeper counterpart.

Relationship between Generalization and Robustness. It is well-known that pruned networks often achieve better generalization than their unpruned counterpart. Table 1 reinforces this claim, in which removing 10% of the filters of ResNet56 provides a generalization gain of 0.60 p.p. Furthermore, a pruned network yielded by the (k + 1)-th iteration of pruning might achieve higher generalization than its preceding (i.e., *k*-th) version. According to Hendrycks et al. [16] and Pang et al. [35], convolutional networks with higher generalization incur better robustness. Therefore, it is intuitive to assume that there exists a relationship between generalization and robustness. In this experiment, we show that such a relationship exhibits a weak correlation from the perspective of pruning.

To confirm the aforementioned statement, we ran several iterations of pruning and plotted the relationship between generation and robustness of each pruned network in Figure 2. It is evident from this figure that there is a weak relationship between generalization and robustness. To further validate this, we compute the Pearson's correlation coefficient (r) and notice that the coefficient is below 0.5 on

different adversarial attacks. In particular, the highest correlation coefficient was 0.31, which indicates a weak correlation. The results from Table 1 reinforce these observations, however, from the viewpoint of the structure removed.

These results show that pruned networks that obtain high generalization are not always the best candidates for achieving better robustness. Fortunately, we observe that one-shot pruning (k = 1 in Algorithm 1) provides a network capable of achieving competitive adversarial robustness. Since k iterations of pruning mean k stages of fine-tuning (see Algorithm 1), one practical benefit of one-shot pruning is that we can avoid multiple epochs of fine-tuning.

**Influence of the Pruning Criterion.** One key ingredient during the pruning process is the criterion for assigning structure importance. According to previous works [14, 23, 41], the pruning criterion plays a role in preserving accuracy after pruning, as it decides which structures are irrelevant/redundant to the network. Depending on the strategy for measuring importance, the pruning criterion can also become a bottleneck in the computational cost of the pruning process [32]. In this experiment, we investigate the behavior of the pruning criterion on adversarial robustness.

Table 2 shows the robustness of (pruned) ResNet56 when the pruning process employs different criteria for deciding which filters to remove. On the semantic-preserving and FGSM attacks, most pruning criteria provided similar gains



Figure 2. Relationship between generalization and robustness of pruned networks. Each point within the plot represents a pruned network yielded throughout the iterations of pruning; thus, k points mean k iterations of pruning. On different attacks, generalization and robustness are uncorrelated, implying that pruned networks with higher generalization do not incur better robustness.

Table 2. Gains in robustness and generalization provided by removing structures of ResNet56 based on different pruning criteria. Differently from the standard assessment metric – accuracy on clean images – evaluating the pruning criteria on adversarial images reveals a shift in their effectiveness. In other words, top-performance criteria for clean images might perform poorly on adversarial images.

Pruning Criterion	Semantic	Occlusion	FGSM	Clean	Average
$\ell_1$ -norm	(+) <b>1.64</b>	(-) 0.80	(+) 3.75	(+) 0.38	(+) 1.24
ExpectedABS [41]	(+) 0.96	(-) 0.09	(+) <b>4.29</b>	(+) 0.51	(+) 1.41
HRank [29]	(+) 0.93	(+) <b>2.92</b>	(+) 3.18	(+) 0.39	(+) 1.85
KlDivergence [32]	(+) 0.82	(+) 0.73	(+) 3.00	(+) 0.34	(+) 1.22
PLS [23]	(+) 1.58	(+) 2.76	(+) 3.68	(+) <b>0.60</b>	(+) 2.15

in robustness. However, on the occlusion attack, the  $\ell_1$ -norm and the criterion by Tan and Motani [41] performed poorly. We shall see later that other pruning settings also affect the gains in the occlusion attack.

According to the average improvement (last column in Table 2), which includes gains in clean images, the best criterion was PLS. In addition to these achievements, PLS exhibits other benefits, e.g., low memory consumption and execution time, which are attractive properties for resource-constrained scenarios. Such advantages are the reason we focus on PLS throughout our study. We highlight that all the pruning criteria evaluated are capable of improving robustness, except  $\ell_1$ -norm and ExpectedABS on the occlusion attack. Thus, our main finding – removing structures from convolutional networks improves their robustness – is unbiased by a specific combination of pruning parameters.

Interestingly, the results in Table 2 suggest exploring robustness as an additional metric for evaluating the pruning criteria. For example, considering clean images only, the difference between the best and worst criteria is within one p.p, which is assumed a negligible difference [3]. On the other hand, this difference increases when we consider the occlusion and FGSM attacks. Additionally, criteria performing well on some attacks have their results notably decreased on others (e.g.,  $\ell_1$ -norm and KIDivergence [32]), thus demonstrating an unexpected instability. Thereby, we could employ robustness as an additional metric for evaluating the effectiveness of the pruning criteria, complementing previous efforts on improving the pruning evaluation [3].

**Effect of Pruning a Single Element.** In all the experiments so far, we analyze the effect of pruning on adversarial images when we remove a percentage (i.e., a small subset) of filters or layers followed by fine-tuning. In this experiment, we demonstrate that even removing a single element (i.e., only one filter) of a network influences their robustness.

To confirm the claim above, we perform the following process. First, we define E as a set containing all possible<sup>2</sup> filters or layers we could eliminate. Afterwards, we remove each element  $e \in E$  individually and measure the

generalization of the resulting network  $\mathcal{F}'$ . It is important to emphasize that  $\mathcal{F}'$  corresponds to  $\mathcal{F}$  without one element e. Then, we add the elements whose removal provides a loss in generalization of  $\{0, 1, 2, 3, 4, 5\}$  percentage points into a corresponding set  $L_{i,i \in \{0,1,2,3,4,5\}}$ . Additionally, these elements are added into some  $L_i$  only if they provide a gain in robustness. For example, we insert the elements  $e \in E$  that preserve the generalization (i.e., zero loss) into  $L_0$  and the elements that drop generalization in up to one p.p into  $L_1$ . Roughly speaking, these steps allow us to evaluate our problem statement under different values of generalization and when the pruning process removes only one element. Algorithm 2 summarizes these steps. Note that we prevent any adjustment in the network parameters after removing an element. Such a setting relies on the fact that removing single filters or layers marginally affects generalization [45, 26, 32] and enables Algorithm 2 to be computationally efficient.

Figure 3 shows the relationship between the degradation in generalization and the percentage of elements  $(\frac{|L_i|}{|S|})$  that, when removed individually, achieve better adversarial robustness than the unpruned network. Taking into account filters as elements in Algorithm 2 and zero drop in generalization ( $L_0$ ), we observe that less than 10% of the filters yield a pruned network with gains in robustness. This percentage substantially increases when we assume a drop in generalization of up to one p.p, for which it increased from 3.07% to 39.88% and from 8.04% to 80.36% on the occlu-

Algorithm 2 Pruning single Elements
<b>Input:</b> Unpruned Network $\mathcal{F}$
<b>Input:</b> Elements to Prune $S$ (filters or layers)
for $i \in \{0, 1, 2, 3, 4, 5\}$ do
$L_i \leftarrow \emptyset$
$\sigma \leftarrow Acc_{clean}(\mathcal{F}) - i \triangleright \text{Degradation in Generalization}$
for $e \in E$ do
$\mathcal{F}' \leftarrow \mathcal{F} \setminus e \triangleright \operatorname{Removes} e \text{ from } \mathcal{F}$
if $Acc_{adv}(\mathcal{F}') > Acc_{adv}(\mathcal{F}) \land Acc_{clean}(\mathcal{F}') > \sigma$ then
$L_i \leftarrow L_i \cup \{e\}$
end if
end for
end for

<sup>&</sup>lt;sup>2</sup>We highlight that some structures must remain unchanged (unpruned) due to technical details of the pruning process such as mismatching between feature maps.



Figure 3. Percentage of structures (y-axis) that, when removed individually, increase robustness given a loss in generalization (xaxis). Take the FGSM attack as an example, with a negligible drop in generalization of one percentage point (dashed line) 80%of the filters, if removed individually, are capable of achieving better adversarial robustness than the unpruned network. Importantly, these results consider removing a structure without fine-tuning.

sion and FSGM attacks, respectively. The curves rapidly saturate on higher values of degradation, indicating that no more filters can yield a pruned network with gains in robustness, even increasing degradation in generalization.

We observe a different behavior when considering layers as elements in Algorithm 2, Figure 3 (right). For example, when removed individually, there was no layer capable of preserving generalization while improving adversarial robustness (it says  $L_0$  is empty). Even evaluating the drop in 1 p.p, the percentage of layers that improve robustness is below 20% on all attacks.

The previous analyses demonstrate that removing single elements without adjusting the network parameters enables us to obtain a certain degree of robustness. We can increase the number of elements satisfying these conditions by controlling the degradation in generalization.

Training from Scratch vs. Fine-tuning. A broad discussion in pruning is how to adjust the network parameters after removing its structures. This adjustment is responsible for recovering the network's predictive ability and plays an important role in the computational cost of the pruning process. In this line of research, previous works either training the pruned network from scratch [30, 54] or fine-tuning it for some epochs [26, 31]. Regarding the first scheme, there are two distinct ways of assigning the network parameters: (i) randomly initialize the parameters [30] and (ii) inherit the random initialization of the unpruned network – winning tickets [8]. In this experiment, we analyze the impact of these schemes on adversarial robustness. To training from scratch with parameters randomly initialized, following the setup by Liu et al. [30], we train a pruned network for the same number of epochs as its unpruned version (Scratch-B) and also consider doubling the epochs (Scratch-E). To the winning-ticket (W-ticket for short) scheme, we inherit the initialization from the unpruned network at epoch 50. We emphasize that all these schemes take into account only clean images.

Table 3. Adversarial robustness considering different schemes to adjust the parameters of the pruned network. Scratch-E, Scratch-B and winning-ticket (W-ticket) train the pruned network from scratch (see the text for details) while fine-tuning adjusts the current parameters of the pruned network for a few epochs. Except for W-ticket, all schemes provide positive average improvements, with fining-tuning achieving the best gains.

	0 0	U		
	Semantic	Occlusion	FGSM	Avg.
Scratch-E	(+) 1.72	(-) 0.60	(+) 1.64	(+) 0.92
Scratch-B	(+)1.25	(-) 0.71	(+) 3.64	(+) 1.39
W-ticket	(+) 1.13	(-) 5.41	(+) 1.40	(-) 0.96
Fine-Tuning	(+) 1.58	(+) <b>2.76</b>	(+) 3.68	(+) <b>2.67</b>

Table 3 shows the gains in the robustness of ResNet56 using different schemes to adjust its parameters after the pruning process. In general, fine-tuning provided better improvements than the training from scratch schemes (Scratch-B/E and W-ticket). Interestingly, none of the training from scratch schemes was capable of providing gains on the occlusion attack. Fine-tuning, on the other hand, provided an improvement of 2.76 p.p. Overall, fine-tuning provided an average improvement of 2.67 p.p., outperforming the best training from scratch (Scratch-B) by 1.28 p.p.

While Table 3 shows gains in robustness, we observe the same trend on generalization, where fine-tuning outperformed Scratch-B/E and W-ticket by 0.28, 0.69 and 0.01 p.p, respectively. These experiments reinforce previous evidence that fine-tuning often leads to better generalization than training from scratch [52]. Our analysis complements this hint, but for the robustness yielded by pruning.

Adversarial Transferability. In previous experiments, we evaluate the adversarial robustness against FGSM using its standard setting – a network is attacked by adversarial images crafted from its parameters. Here, we relax this assumption and assume that the FGSM attack is unaware of the victim (target) network. To evaluate this transferability of adversarial images, we craft adversarial images from the parameters of a source network and attack a target network. It is worth mentioning that the semantic-preserving and occlusion attacks are model-agnostic (i.e., they craft adversarial images without access to the network details). Therefore, such attacks were not used in this experiment.

Table 4 shows the results of the FGSM attack on different target and source architectures. From these results, when the source and target architectures are the same (diagonal – the attack knows the victim), the adversarial robustness is below 30%. However, when the attack lacks access to the target architecture (off-diagonal), the adversarial robustness exhibited substantial gains. Interestingly, the adversarial robustness is even better when the source and target architectures are completely different, for example, the source is MobileNetV2 (including its pruned version) and the target is ResNet56. In addition, except when Mo-

Table 4. Adversarial robustness when a model-specific attack (FGSM) yields adversarial images with different architectures including their pruned versions. Source and target indicate the network (to be precise its parameters) used to craft and evaluate the adversarial images, in this order. Thus, the off-diagonal shows robustness when the attack is unaware of the victim architecture or the defense mechanism (e.g., pruning). Bold values indicate the best robustness. Except when MobileNetV2 is the source, pruning an architecture provides the best adversarial robustness. These results are even better when the source and target architectures are different.

Target	ResNet56	ResNet56 + Pruning	MobileNetV2	MobileNetV2+Pruning
ResNet56	21.60	32.73	52.55	53.25
ResNet56 + Pruning	33.91	25.28	53.67	53.68
MobileNetV2	70.28	70.25	24.21	27.17
MobileNetV2 + Pruning	69.69	70.01	28.29	24.84

bileNetV2 is the source, the pruned versions of the target architectures achieved even better results. Altogether, these results reinforce that pruned networks become more robust to adversarial images. Such evidence persists even when these images come from different architectures.

**Comparison with Competing Defense Mechanisms.** In this experiment, we compare the effectiveness of pruning with state-of-the-art defense mechanisms against adversarial images. It is worth mentioning that our focus is on showing the potential of pruning as a novel defense mechanism but not on pushing the state of the art in adversarial robustness. Following previous works [6, 27], we evaluate several defense mechanisms on the ImageNet-C dataset [16].

Table 5 shows the gains (or degradation) in robustness and generalization of ResNet50 when we employ different defense mechanisms. According to Table 5, the data augmentation by style transfer (Stylized for short) and MixUp exhibited degradation in the average improvement. Thus, these defense mechanisms fail to satisfy the dilemma between generalization and robustness. On the other hand, pruning provides gains in both metrics. Specifically, when the pruning process considers removing filters, the average improvement is 11.38 and 3.97 p.p higher than MixUp and Stylized, respectively. These gains are slightly smaller

Table 5. Comparison of pruning with state-of-the-art defense mechanisms. Compared with competing defense mechanisms, pruning obtained one of the best average improvements. Such results indicate that pruning is capable of satisfying the dilemma between robustness and generalization. The symbols  $\dagger$  and  $\ddagger$  indicate results taken from Li et al. [27] and Chun et al. [6].

Defense	Robustness	Generalization	Average
Stylized †	(-) 2.29	(-) 16.20	(-) 9.24
MixUp ‡	(-) 4.77	(+) 1.10	(-) 1.83
Cutout ‡	(+) 1.39	(+) 0.75	(+) 1.06
CutMix ‡	(+) 1.71	(+) 2.07	(+) 1.89
Shape-Texture †	(+) 7.50	(+) 0.50	(+) 4.00
Pruning Filters	(+) 1.14	(+) 3.15	(+) 2.14
Pruning Layers	(+) 1.20	(+) 3.03	(+) 2.11

when the pruning process removes layers. Compared to the defense mechanisms that simultaneously improve robustness and generalization (e.g., Cutout and CutMix), pruning achieved one of the best average improvements. In particular, only the sophisticated data augmentation by Li et al [27] achieved a better average improvement than pruning.

It is worth mentioning that the mechanisms in Table 5 require training the network from scratch. Pruning, however, obtained competitive results with only some epochs (i.e., 12 epochs) of fine-tuning. This advantage becomes pruning an attractive defense mechanism, mainly to off-the-shelf networks and scenarios with a limited training budget. Finally, since our pruning process employed only clean images, it is orthogonal to defense mechanisms based on data augmentation. Therefore, we could combine pruning with these mechanisms, yielding more robust and efficient networks.

#### **5.** Conclusions

We investigate the behavior of pruning through the lens of adversarial robustness. We empirically show that pruning filters and/or layers of convolutional networks increase their adversarial robustness. Such evidence demonstrates that pruned networks do not inherit the adversarial vulnerability of their over-parametrized (unpruned) counterpart. Furthermore, pruning preserves generalization; thus, it efficiently satisfies the dilemma between robustness and generalization. We confirm these findings considering only clean images during the pruning process. One practical benefit of this setup is to avoid adversarial training and, hence, its drawbacks. Additionally, our analysis enables us to design an effective defense mechanism that ignores the settings and assumptions of the attack. Compared to state-of-the-art defense mechanisms, pruning obtained one of the best tradeoffs between robustness and generalization.

#### Acknowledgments

The authors would like to thank FAPESP (grant #2017/12646-3) and CNPq (grant #309330/2018-1).

## References

- Yang Bai, Yuyuan Zeng, Yong Jiang, Shu-Tao Xia, Xingjun Ma, and Yisen Wang. Improving adversarial robustness via channel-wise activation suppressing. In *International Conference on Learning Representations (ICLR)*, 2021. 3
- [2] Brian Bartoldson, Ari S. Morcos, Adrian Barbu, and Gordon Erlebacher. The generalization-stability tradeoff in neural network pruning. In *Neural Information Processing Systems (NeurIPS)*, 2020. 2, 4
- [3] Davis W. Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John V. Guttag. What is the state of neural network pruning? In *Conference on Machine Learning and Systems (MLSys)*, 2020. 6
- [4] Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C. Duchi, and Percy Liang. Unlabeled data improves adversarial robustness. In *Neural Information Processing Systems (NeurIPS)*, 2019. 1
- [5] Ting-Wu Chin, Ruizhou Ding, Cha Zhang, and Diana Marculescu. Towards efficient model compression via learned global ranking. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 4
- [6] Sanghyuk Chun, Seong Joon Oh, Sangdoo Yun, Dongyoon Han, Junsuk Choe, and Youngjoon Yoo. An empirical evaluation on robustness and uncertainty of regularization methods. In *International Conference on Machine Learning (ICML) Workshop on Uncertainty and Robustness in Deep Learning*, 2019. 1, 2, 3, 4, 8
- [7] Angela Fan, Edouard Grave, and Armand Joulin. Reducing transformer depth on demand with structured dropout. In *International Conference on Learning Representations (ICLR)*, 2020. 2, 4
- [8] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Rep*resentations (ICLR), 2019. 7
- [9] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations (ICLR)*, 2019. 3
- [10] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2015. 3
- [11] Jinyang Guo, Wanli Ouyang, and Dong Xu. Multidimensional pruning: A unified framework for model compression. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 4

- [12] Kai Han, Yunhe Wang, Qiulin Zhang, Wei Zhang, Chunjing XU, and Tong Zhang. Model rubiks cube: Twisting resolution, depth and width for tinynets. In *Neural Information Processing Systems (NeurIPS)*, 2020. 4
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Conference on Computer Vision and Pattern Recognition (CVPR), 2016. 4
- [14] Yang He, Yuhang Ding, Ping Liu, Linchao Zhu, Hanwang Zhang, and Yi Yang. Learning filter pruning criteria for deep convolutional neural networks acceleration. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5
- [15] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2019. 5
- [16] Dan Hendrycks and Thomas G. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations (ICLR)*, 2019. 1, 3, 4, 5, 8
- [17] Dan Hendrycks, Kimin Lee, and Mantas Mazeika. Using pre-training can improve model robustness and uncertainty. In *International Conference on Machine Learning (ICML)*, 2019. 1
- [18] Dan Hendrycks, Norman Mu, Ekin Dogus Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. In *International Conference on Learning Representations* (*ICLR*), 2020. 1, 3
- [19] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1
- [20] Ting-Kuei Hu, Tianlong Chen, Haotao Wang, and Zhangyang Wang. Triple wins: Boosting accuracy, robustness and efficiency together by enabling inputadaptive inference. In *International Conference on Learning Representations (ICLR)*, 2020. 1
- [21] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q. Weinberger. Deep networks with stochastic depth. In *European Conference on Computer Vision* (ECCV), 2016. 2, 4
- [22] Zehao Huang and Naiyan Wang. Data-driven sparse structure selection for deep neural networks. In *European Conference on Computer Vision (ECCV)*, 2018.
  4

- [23] Artur Jordao, Fernando Yamada, and William Robson Schwartz. Deep network compression based on partial least squares. *Neurocomputing*, 406:234–243, 2020.
  4, 5, 6
- [24] Yigitcan Kaya, Sanghyun Hong, and Tudor Dumitras. Shallow-deep networks: Understanding and mitigating network overthinking. In *International Conference on Machine Learning (ICML)*, 2019. 1
- [25] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *European Conference on Computer Vision (ECCV)*, 2020. 1
- [26] Jiashi Li, Qi Qi, Jingyu Wang, Ce Ge, Yujian Li, Zhangzhang Yue, and Haifeng Sun. OICSR: out-inchannel sparsity regularization for compact deep neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 4, 6, 7
- [27] Yingwei Li, Qihang Yu, Mingxing Tan, Jieru Mei, Peng Tang, Wei Shen, Alan L. Yuille, and Cihang Xie. Shape-texture debiased neural network training. In *International Conference on Learning Representations* (*ICLR*), 2021. 1, 8
- [28] Yingwei Li, Qihang Yu, Mingxing Tan, Jieru Mei, Peng Tang, Wei Shen, Alan L. Yuille, and Cihang Xie. Shape-texture debiased neural network training. In *International Conference on Learning Representations* (*ICLR*), 2021. 3
- [29] Mingbao Lin, Rongrong Ji, Yan Wang, Yichen Zhang, Baochang Zhang, Yonghong Tian, and Ling Shao. Hrank: Filter pruning using high-rank feature map. In Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 4, 6
- [30] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. In *International Conference on Learning Representations (ICLR)*, 2019. 7
- [31] Jian-Hao Luo, Hao Zhang, Hong-Yu Zhou, Chen-Wei Xie, Jianxin Wu, and Weiyao Lin. Thinet: Pruning CNN filters for a thinner net. *Transactions on Pattern Analysis and Machine Intelligence*, 41:2525– 2538, 2019. 4, 7
- [32] Jian-Hao Luo and Jianxin Wu. Neural network pruning with residual-connections and limited-data. In Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 4, 5, 6
- [33] 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, 3

- [34] Chaithanya Kumar Mummadi, Ranjitha Subramaniam, Robin Hutmacher, Julien Vitay, Volker Fischer, and Jan Hendrik Metzen. Does enhanced shape bias improve neural network robustness to common corruptions? In *International Conference on Learning Representations (ICLR)*, 2021. 3
- [35] Tianyu Pang, Xiao Yang, Yinpeng Dong, Hang Su, and Jun Zhu. Bag of tricks for adversarial training. In *International Conference on Learning Representations (ICLR)*, 2021. 3, 5
- [36] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John C. Duchi, and Percy Liang. Adversarial training can hurt generalization. In *International Conference* on Machine Learning (ICML) Workshop on Identifying and Understanding Deep Learning Phenomena, 2019. 1
- [37] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John C. Duchi, and Percy Liang. Understanding and mitigating the tradeoff between robustness and accuracy. In *International Conference on Machine Learning (ICML)*, 2020. 1, 3
- [38] Shahbaz Rezaei and Xin Liu. A target-agnostic attack on deep models: Exploiting security vulnerabilities of transfer learning. In *International Conference* on Learning Representations (ICLR), 2020. 3
- [39] Ali Shafahi, Parsa Saadatpanah, Chen Zhu, Amin Ghiasi, Christoph Studer, David W. Jacobs, and Tom Goldstein. Adversarially robust transfer learning. In *International Conference on Learning Representations (ICLR)*, 2020. 1
- [40] Baifeng Shi, Dinghuai Zhang, Qi Dai, Zhanxing Zhu, Yadong Mu, and Jingdong Wang. Informative dropout for robust representation learning: A shape-bias perspective. In *International Conference on Machine Learning (ICML)*, 2020. 3
- [41] Chong Min John Tan and Mehul Motani. Dropnet: Reducing neural network complexity via iterative pruning. In *International Conference on International Conference on Machine Learning (ICML)*. 2020. 4, 5, 6
- [42] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on International Conference on Machine Learning (ICML)*, 2019. 1, 4
- [43] Brandon Tran, Jerry Li, and Aleksander Madry. Spectral signatures in backdoor attacks. In *Neural Information Processing Systems (NeurIPS)*, 2018. 1
- [44] Andreas Veit and Serge J. Belongie. Convolutional networks with adaptive inference graphs. *International Journal of Computer Vision (IJCV)*, 128:730– 741, 2020. 4

- [45] Andreas Veit, Michael J. Wilber, and Serge J. Belongie. Residual networks behave like ensembles of relatively shallow networks. In *Neural Information Processing Systems (NeurIPS)*, 2016. 4, 6
- [46] 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. In *Neural Information Processing Systems (NeurIPS)*, 2020. 1, 3
- [47] Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E. Gonzalez. Skipnet: Learning dynamic routing in convolutional networks. In *European Conference on Computer Vision (ECCV)*, 2018. 4, 5
- [48] Zuxuan Wu, Tushar Nagarajan, Abhishek Kumar, Steven Rennie, Larry S. Davis, Kristen Grauman, and Rogério Schmidt Feris. Blockdrop: Dynamic inference paths in residual networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 4
- [49] Cihang Xie, Mingxing Tan, Boqing Gong, Jiang Wang, Alan L. Yuille, and Quoc V. Le. Adversarial examples improve image recognition. In *Conference* on Computer Vision and Pattern Recognition (CVPR), 2020. 1, 2, 3, 5
- [50] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan L. Yuille. Mitigating adversarial effects through randomization. In *International Conference* on Learning Representations (ICLR), 2018. 4
- [51] Cihang Xie and Alan L. Yuille. Intriguing properties of adversarial training at scale. In *International Conference on Learning Representations (ICLR)*, 2020. 1, 2, 3, 5
- [52] Mao Ye, Chengyue Gong, Lizhen Nie, Denny Zhou, Adam Klivans, and Qiang Liu. Good subnetworks provably exist: Pruning via greedy forward selection. In *International Conference on Machine Learning (ICML)*, 2020. 7
- [53] Shaokai Ye, Xue Lin, Kaidi Xu, Sijia Liu, Hao Cheng, Jan-Henrik Lambrechts, Huan Zhang, Aojun Zhou, Kaisheng Ma, and Yanzhi Wang. Adversarial robustness vs. model compression, or both? In *International Conference on Computer Vision (ICCV)*, 2019. 4
- [54] Haoran You, Chaojian Li, Pengfei Xu, Yonggan Fu, Yue Wang, Xiaohan Chen, Yingyan Lin, Zhangyang Wang, and Richard G. Baraniuk. Drawing early-bird tickets: Toward more efficient training of deep networks. In *International Conference on Learning Representations (ICLR)*, 2020. 7
- [55] Sangdoo Yun, Dongyoon Han, Sanghyuk Chun, Seong Joon Oh, Youngjoon Yoo, and Junsuk Choe.

Cutmix: Regularization strategy to train strong classifiers with localizable features. In *International Conference on Computer Vision (ICCV)*, 2019. 3, 4

- [56] Chiyuan Zhang, Samy Bengio, and Yoram Singer. Are all layers created equal? In International Conference on Machine Learning (ICML) Workshop on Identifying and Understanding Deep Learning Phenomena, 2019. 4
- [57] Linjun Zhang, Zhun Deng, Kenji Kawaguchi, Amirata Ghorbani, and James Y. Zou. How does mixup help with robustness and generalization? In *International Conference on Learning Representations* (*ICLR*), 2021. 3
- [58] Wangchunshu Zhou, Canwen Xu, Tao Ge, Julian J. McAuley, Ke Xu, and Furu Wei. BERT loses patience: Fast and robust inference with early exit. In *Neural Information Processing Systems (NeurIPS)*, 2020. 1