

# Defense-friendly Images in Adversarial Attacks: Dataset and Metrics for Perturbation Difficulty

Camilo Pestana

camilo.pestanacardeno@research.uwa.edu.au

Wei Liu

wei.liu@uwa.edu.au

David Glance

david.glance@uwa.edu.au

Ajmal Mian

ajmal.mian@uwa.edu.au

The University of Western Australia  
35 Stirling Hwy, Crawley WA 6009, Australia

## Abstract

*Dataset bias is a problem in adversarial machine learning, especially in the evaluation of defenses. An adversarial attack or defense algorithm may show better results on the reported dataset than can be replicated on other datasets. Even when two algorithms are compared, their relative performance can vary depending on the dataset. Deep learning offers state-of-the-art solutions for image recognition, but deep models are vulnerable even to small perturbations. Research in this area focuses primarily on adversarial attacks and defense algorithms. In this paper, we report for the first time, a class of robust images that are both resilient to attacks and that recover better than random images under adversarial attacks using simple defense techniques. Thus, a test dataset with a high proportion of robust images gives a misleading impression about the performance of an adversarial attack or defense. We propose three metrics to determine the proportion of robust images in a dataset and provide scoring to determine the dataset bias. We also provide an ImageNet-R dataset of 15000+ robust images to facilitate further research on this intriguing phenomenon of image strength under attack. Our dataset, combined with the proposed metrics, is valuable for unbiased benchmarking of adversarial attack and defense algorithms.*

## 1. Introduction

Convolutional Neural Networks (CNNs) are used in a wide range of computer vision tasks [20] such as object detection and image classification. However, in 2014, Szegedy et al. [36] demonstrated that small perturbations in an input image can make a CNN model misclassify it. These perturbed images are called Adversarial Examples. As attacks are getting more sophisticated [30], there

is an increasing need for more robust models and better defenses [10].

Pestana et al. [28] suggested that depending on the subset of images selected from the ImageNet [20] validation set, results reported by different adversarial defenses (for ImageNet-based trained models) might differ remarkably. This is a significant problem that tarnishes the validity of current evaluations reported in Adversarial Defense research. To date, there are three main benchmark datasets that are used to evaluate the performance of adversarial attacks and defenses in images, namely MNIST [21], CIFAR10 [19], and ImageNet [20]. Of these datasets, ImageNet is the most complex given the number of classes (1,000) and the input image size (some models resize the images from ImageNet to  $300 \times 300$  pixels or more commonly  $224 \times 224$ ). MNIST and CIFAR10 datasets are commonly used to quickly demonstrate multiple attacks and defense techniques. However, these attacks do not transfer well to other datasets. For example, the defense technique JumpReLU [10] works well on CIFAR10 and MNIST but is not effective on ImageNet.

Hendrycks et al. [17] curated two datasets for ImageNet models that contain 7,500 hard-to-classify images or “Naturally Adversarial Images” and named this dataset as ImageNet-A (where ‘A’ stands for adversarial). Their aim was to provide a new way to measure *model* robustness and uncertainty.

On the other side of the spectrum, there exists a subset of images that is more difficult to perturb, i.e. “Naturally Robust Images”. This was hinted to by Pestana et al. [28]. However, there is no readily available metrics to measure the *data* robustness.

Our finding suggests that images that are more difficult to perturb are also easier to recover when using defenses, even if the defenses are weak. In this paper, we use the term

“defense-friendly” to describe the subset of images that are not only more robust but also recover more easily from an attack using defenses. To further demonstrate the existence of such a group of images, we curated a dataset of defense-friendly images.

The main contributions of our work are:

(1) **Robust Dataset:** Existing datasets are not sufficiently challenging for evaluating adversarial attack/defense mechanisms. In this paper, a method for identifying robust images is proposed and used to construct such a dataset.

**For benchmarking:** We curated a dataset (ImageNet-R) containing easy to classify, robust to perturbation and defense-friendly images. This dataset will be released publicly to augment the ImageNet-A dataset. We use this dataset to validate the aforementioned claims about the shortcomings of existing datasets.

**For Robustness Research:** Our dataset ImageNet-R can inspire future research regarding the phenomenon of robust/defense-friendly images and create new lines of research.

**For Adversarial Research:** A dataset with a large proportion of robust images might give the wrong impression that the performance of a defense is better than it actually is. On the other hand, using a robust (hard-to-perturb) dataset might provide insights on how to improve existing attacks.

(2) **Robustness Metrics:** It is important to quantify the robustness of existing datasets, we therefore propose three metrics, namely ARD, AMP, and ADF score to measure the characteristics of a dataset with respect to adversarial attacks and defense algorithms. Given that those metrics are expensive to compute, a proxy method based on machine learning (i.e., predicting the robustness of an image) is introduced. We show the effectiveness of this proxy as a useful tool for creating unbiased benchmarks.

## 2. Problem Definition

### 2.1. Adversarial Attacks and Defenses

Let  $x \in \mathbb{R}^{H \times W \times 3}$  denote the original image without perturbations, where  $H$  is height,  $W$  is width and there are three color channels, usually (R)ed, (G)reen and (B)lue. Let  $y \in \mathbb{R}^n$  be the probabilities of the predicted labels. Given an image classifier  $C : x \rightarrow \{1, 2, \dots, n\}$ , e.g., for an ImageNet dataset [7]  $n = 1,000$ , an untargeted attack aims to add a perturbation  $p$  to  $x$  to compute  $x_{adv}$ , such that  $C(x) \neq C(x_{adv})$ , where  $p$  could be a grayscale image of the same size  $p \in \mathbb{R}^{H \times W}$ ; or a colored image of the same size  $p \in \mathbb{R}^{H \times W \times 3}$ . The calculation  $x_{adv} = x + p$  is constrained such that the perturbation in  $x_{adv}$  is imperceptible for the human eye, e.g.,  $d(x, x_{adv}) \leq \epsilon$  for a distance function  $d()$  and a small value  $\epsilon$ . In the context of adversarial attacks, the distance metric  $d()$  is the  $L_p$  norm of the difference between the original image  $x$  and the adversarial image  $x_{adv}$ . In this paper, we consider attacks that optimise

the  $L_2$  and  $L_\infty$  norms only. If the computed  $x_{adv}$  satisfies the condition  $C(x) \neq C(x_{adv})$  under the given set of constraints, the attack is considered successful. In addition, let  $D()$  denote a defense function; if  $C(x) \neq C(x_{adv})$  then  $D()$  should ideally behave such that  $C(D(x_{adv})) = C(x)$ .

### 2.2. Easy, Robust and Defense-friendly Images

Given a classifier  $C_i()$ , an image  $x$  is considered *easy* if different classifiers trained on the same data, e.g., DenseNet121 [18], Inception-v3 [37], Vgg16 [33], agree with the prediction such that  $C_0(x) = C_1(x) = \dots = C_n(x)$ . An image  $x$  is considered  $\epsilon$ -robust if its adversarial version  $x_{adv}$  is classified correctly such that  $C(x_{adv}) = C(x)$  for a small  $\epsilon$  value. Therefore, an easy image can be considered a special case of the most general definition of robust images where  $\epsilon = 0$ . In our experiments, we found that  $\epsilon$ -robust images are defense friendly at higher  $\epsilon$ , where  $\epsilon$  is the highest perturbation that an image can withstand without misclassification. Moreover, at the same perturbation level,  $\epsilon$ -robust images and epsilon defense friendly images are a complementary set, in that they do not have intersections. Therefore, images that are robust have a higher chance of recovering accuracy with correct class prediction at inference time. In other words, robust images are often “defense-friendly”. As shown in Table 1, a test set with more robust images return better accuracy under the same defense mechanism.

## 3. Related Work

We briefly discuss popular adversarial attacks and defenses proposed in the literature that will be used in our experiments. We only focus on adversarial examples in the domain of image classification. In addition, techniques to obtain statistical features from images are discussed along with interpretability methods in Deep Learning.

### 3.1. Attack Algorithms

Gradient-based attacks are more powerful than Non-gradient-based attacks and usually less computationally expensive [2], hence a defense against them is practically more meaningful. We use gradient-based attacks to test and compare our robust datasets and test the performance of some defenses using different sets of images.

**Fast Gradient Signed Method (FGSM):** This attack is one of the first computationally efficient single step adversarial attacks mentioned in the literature and introduced by Goodfellow et al. [12]. FGSM is calculated like  $x_{adv} = x + \epsilon * \text{sign}(\nabla_x J(\theta, x, y))$  where  $x_{adv}$  is the adversarial image and  $x$  is the original image. In the cost function  $J(\theta, x, y)$ ,  $\theta$  represents the network parameters and  $y$  the ground truth label. Moreover,  $\epsilon$  is used to scale the noise and is usually a small number (e.g.  $\epsilon=0.01$ ). A sign function is applied to the gradients of the loss with respect to the

Table 1: Comparison of Non-Robust (NR1) and Robust (R1) datasets accuracy (%) under different attacks and defenses. Images from R1 dataset are easier to defend not only for a single model but for all of them. Even though, the transferability of this property “defense-friendly” is not 100% between models, it shows that compared to a non-robust dataset (NR1), the likelihood of a robust image to recover from the wrong class prediction when using a defense is significantly higher. These robust images are also significantly more difficult to perturb and attacks would need a higher level of perturbation to be more effective. Higher perturbations levels compromise imperceptibility.

Data	Models	No Defense									BaRT Defense									ResUpNet Defense											
		FGSM			PGD			DDN			FGSM			PGD			DDN			FGSM			PGD			DDN					
		Clean	$\epsilon$		$\epsilon$		n		Clean	$\epsilon$		$\epsilon$		n		Clean	$\epsilon$		$\epsilon$		n		Clean	$\epsilon$		$\epsilon$		n			
			0.01	0.02	0.04	0.01	0.02	0.04	20	40	60	Clean	0.01	0.02	0.04	0.01	0.02	0.04	20	40	60	Clean	0.01	0.02	0.04	0.01	0.02	0.04	20	40	60
Non-Robust (NR1)	vgg	68.7	0	0	0	0	0	0	0	0	44.4	1.3	0.6	0	1.1	0.7	0.0	4.0	5.0	4.0	41.4	9.3	6.7	0.1	12.0	9.0	6	31.8	31.8	31.8	
	resnet	81.2	3.1	6.2	3.1	0	0	0	0	0	62.7	4.5	8.7	4.9	5.9	4.1	2.0	8.0	8.0	11.0	61.0	31	28.1	28.1	28.0	14.0	6	47.3	47.3	47.3	
	densenet	87.5	15.6	9.4	6.2	0	0	0	0	0	76.3	4.3	15.2	9.8	13	9.3	8.0	22.0	17.0	16.0	84.0	47	28.4	25.7	37.0	16.0	6	72.6	72.6	72.6	
	inception	85.9	14.4	13.4	9.4	0	0	0	66.6	66.6	66.6	58.1	4.7	2.9	3.5	5.1	2.9	3.7	8.8	8.9	8.4	60.9	19.4	13.1	3.1	9.6	0	0	39.04	39.04	39.04
	mobilenet	85.9	12.8	9.1	7.2	0	0	0	73.4	73.4	73.4	34.1	4.1	2.9	1.2	4.4	1.9	1.7	5.3	6.9	7.0	40.6	22.5	1.1	0.0	15.6	0	0	34.4	34.4	34.4
	shufflenet	87.5	8.1	6.0	4.1	0	0	0	70.3	70.3	70.3	37.2	1.4	1.0	0.5	1.7	1.3	1.2	7.2	5.4	5.9	32.8	29.2	4.7	1.6	12.5	1.6	1.6	23.4	23.4	23.4
mnasnet	90.6	11.2	3.4	4.1	0	0	0	78.1	78.1	78.1	39.1	4.1	3.2	1.7	2.9	2.9	1.7	7.3	7.1	6.8	40.6	28.8	9.4	1.6	20.3	7.8	3.1	34.4	34.4	34.4	
Robust (R1)	vgg	89.8	83.3	82.2	79.5	54.3	26.3	11.2	83.2	83.2	83.2	89.8	76.9	73.7	72.8	55.4	33.8	19.9	81.6	80.1	80.3	61.1	71.5	71.1	70.4	69.2	65.8	59.9	62.8	62.8	62.8
	resnet	98.0	100	100	100	99.8	43.9	17.5	100	100	100	98.3	94.7	95.2	95.1	93.9	57.2	34.9	98.0	98.7	98.6	83.9	87.4	87.3	86.9	86.0	83.9	79.2	83.9	83.9	83.9
	densenet	99.2	100	100	100	100	55.5	22.1	100	100	100	99.2	97.0	97.1	96.9	95.9	67.0	40.9	99.3	98.7	99.0	90.7	92.8	92.8	92.0	90.6	88.4	83.5	90.7	90.7	90.7
	inception	100	97.9	97.9	97.1	87.2	64.5	42.9	98.1	98.1	98.1	98.3	94.0	93.0	94.0	86.8	72.7	52.9	96.7	97.2	96.5	86.9	90.5	90.4	90.3	89.0	87.1	83.3	86.5	86.5	86.5
	mobilenet	99.2	89.0	88.2	85.9	55.4	23.2	8.4	88.8	88.8	88.8	90.2	82.0	84.0	81.0	60.0	34.7	22.3	87.7	86.3	86.5	74.7	76.0	75.6	74.8	73.5	70.8	65.7	74.5	74.5	74.5
	shufflenet	99.1	82.4	80.8	77.6	41.3	16.6	5.8	82.5	82.5	82.5	87.4	76.0	75.0	71.0	48.4	28.7	17.7	80.8	78.6	79.3	65.0	70.2	69.8	69.0	66.7	62.7	56.7	65.8	65.8	65.8
mnasnet	99.6	90.5	89.6	87.7	57.9	23.6	8.6	90.7	90.7	90.7	92.8	81.9	81.8	82.2	61.4	36.1	24.9	88.6	87.3	87.6	71.1	72.6	72.4	71.7	70.2	67.0	61.7	71.0	71.0	71.0	

input image to compute the final perturbation.

**Projected Gradient Descent (PGD):** A stronger iterative version of FGSM is the PGD attack, which is considered one of the strongest attacks and used as a benchmark to measure the robustness of many defenses in the literature [24]. It works similarly to FGSM, however, in this iterative version a small perturbation step  $\alpha$  is applied in every step. The “projection”, for example, for an  $L_2$  norm in a 2 dimensional space would mean to move a point to the closest point inside a circumference where the center is the origin (original image without perturbation).

**Carlini and Wagner (C&W):** Carlini and Wagner [4] introduced three attacks in response to one of the first adversarial defenses in the literature called Defensive Distillation. These attacks are hardly noticeable by humans, given that they constrain the  $L_p$ -norm. The C&W attacks are shown to be highly effective against some types of defenses. A C&W attack that is constrained by the  $L_2$  norm is considered one of the strongest attacks in the literature [31], however, this attack can be computationally expensive, often requiring thousands of iterations to calculate an adversarial image.

**Decoupling Direction and Norm (DDN):** The DDN attack is an efficient method that was introduced by Rony et al. [31]. DDN optimizes the number of iterations needed while achieving state-of-the-art results, even comparable to C&W  $L_2$  constrained attacks.

## 3.2. Defenses

Recently, a number of studies have explored ideas to make Deep Neural Networks (DNN) more robust against adversarial attacks. Many of those defense methods are still unable to achieve true robustness to all adversarial inputs. Currently, the most effective defense strategies modify the DNN training process to improve robustness against adversarial examples [32]. However, they are trained to defend against specific attacks, limiting their real-world applications. In contrast, there is another research line on defenses which aims to be attack and model agnostic. This line pre-processes the images instead of modifying models or applying specific defenses to them. These two different approaches are: defenses implemented in the training process [35, 38, 22] and defenses that are applied as an image denoising operation [29, 40]. Below we describe some of the defenses with respect to these two categories. However, our aim in this paper is to have a model and defense agnostic approach in the experiments. For that reason, our focus is mainly on the image preprocessing or image denoising techniques.

### 3.2.1 Adversarial Training

The aim of adversarial training is to increase the DNN’s robustness by adding adversarial images to the training set [12], [23], [32]. These methods effectively enhance robustness against adversarial attacks, but they lack generali-

sation to unknown attacks [39]. However, as demonstrated by [25] and [26], this type of defense is not effective to non-gradient-based attacks.

### 3.2.2 Image Denoising

The approach of pre-processing input examples to remove adversarial perturbations has the advantage of being model-agnostic. Moreover, it can be combined with any other defense strategies. Bhagoji et al. [3] proposed a Principal Component Analysis (PCA) method on images to reduce their dimensionality and therefore reduce noise. Alternatively, Das et al. [6] proposed leveraging JPEG compression as a pre-processing step for adversarial defenses.

The defense Barrage of Random Transformations (BaRT) combines different input transformation defenses in the literature such as Feature Squeezing [41], Blurring Filters [5] and JPEG Compression [9], [13] into a pipeline. He et al. [16] showed that combining weak defenses does not create a stronger defense. However, Raff et al. [29] demonstrated that randomly selecting transformations from a big pool of transformations into a single barrage (BaRT) is a defense capable of resisting the strongest attacks such as PGD. This defense has a pool of 25 different transformations that can be randomly selected from the pool. The parameters for every transformation are selected randomly as well as the order  $k$  of the number of transformations selected. Then, once they are combined in a single pipeline the transformations are applied sequentially. By evaluating this defense, we are indirectly testing and comparing the performance against 25 different input transformations.

ResUpNet defense [28] exploits the insight that the Y channel from the YCbCr is more relevant in terms of adversarial perturbation and focuses on recovering the Y-channel in that specific color space before converting it back to RGB. This defense is reported to achieve better than the BaRT defense.

### 3.3. Image Features

Convolutional Neural Networks (CNNs) have become a very popular approach to extract features from images. However, for image analysis other statistical techniques also exist such as the grey level co-occurrence matrix [14] to measure the texture of an image. Ideally, in texture analysis the aspects of an image should be rotationally invariant. A common practice is to calculate the grey co-occurrence matrix using different angles (e.g. 0, 45, 90, 135) [34]. Hence, for our experiments we calculate Grey Level Co-occurrence Matrices (GLCM) rotating the image at the 4 different angles mentioned. From those GLCM it is possible to extract the texture properties below:

1) **Contrast** refers to the calculation of the intensity contrast between a pixel and its neighbors i.e.  $\sum_{i,j=0}^{N-1} P_{i,j} (i-j)^2$ .

2) **Dissimilarity** is the variation of grey-level pairs computed as  $\sum_{i,j=0}^{N-1} P_{i,j} |i-j|$ .

3) **Homogeneity** measures the tightness of distribution of the elements in the GLCM to the GLCM diagonal,  $\sum_{i,j=0}^{N-1} \frac{P_{i,j}}{1+(i-j)^2}$ .

4) **Angular Second Moment (ASM)** is a measure of the homogeneity of an image computed as  $\sum_{i,j=0}^{N-1} P_{i,j}^2$ .

5) **Energy** refers to the information related to image homogeneity. It has a low value if the pairs are similar and high values otherwise and is computed as  $\sqrt{\sum_{i,j=0}^{N-1} P_{i,j}^2}$ .

6) **Correlation** measures linear dependency between pixels. It has high values when the pixels are uniformly distributed and is computed as  $\sum_{i,j=0}^{N-1} P_{i,j} \frac{(i-\mu_i)(j-\mu_j)}{\sqrt{(\sigma_i^2)(\sigma_j^2)}}$ .

## 4. Defense-friendly Dataset

While evaluating the performance of adversarial attacks and defenses for ImageNet, it is a common practice to take a subset of images from ImageNet to test the performance of different techniques (attacks or defenses). However, some subsets of images perform much better or worse than a randomly selected subset [28]. This can lead to incorrect assumptions or conclusions with regards to the performance of models or defenses implemented. In this section, we explain the process we used to collect a set of images, outside of the ImageNet validation dataset, that are more resilient to attacks than randomly selected images and recover easily when using a defense (defense-friendly images). The source code and datasets will be made publicly available at <https://github.com/elcronos/Defense-Friendly>

### 4.1. Implementation Details

We used the ResUpNet [28] and BaRT [29] in the PyTorch [27] library. In addition, AdverTorch [8] was used to create adversarial images for three different classifiers Vgg16 [33], ResNet50 [15], and DenseNet121 [18]. Those models have unique architecture mechanisms and differ in the number of layers, which makes them a good benchmark to test the defenses using different datasets. For  $L_\infty$  constrained attacks such as FGSM and PGD, we use  $\epsilon$  values of 0.01, 0.02 and 0.04 for images with values in the range [0,1]. In the case of DDN attack, instead of  $\epsilon$ , we use the number of iterations  $n$  as the hyper-parameter, with values 20, 40, and 60.

### 4.2. Curating a Defense-friendly dataset

To demonstrate the existence of defense-friendly images and to compare their performance with randomly selected images, we curated a dataset of defense-friendly images. For this experiment, we used pre-trained ImageNet classifiers. We decided to create a new dataset using Flickr to prevent any data leakage between the training and test datasets.

In machine learning data leakage occurs because of the use of information in the model training process which is not available at prediction time. As a consequence, the predictive scores are overestimated. We assume that very few or none of the images we collected are included in either the training ImageNet dataset or test dataset. Since ImageNet was collected scrapping the web in 2012, and we only collected images from the last two years from the Flickr API, we assume that the probability of having repeated images remains extremely low.

Flickr is a website that has millions of images shared by users. We used the Flickr API to download 1,000 images per each ImageNet class. The API uses a keyword system to download the images. The keywords used to download the images will also be shared in our GitHub repository.

We used DenseNet-121, ResNet-50 and Inception-v3 pretrained models to further filter these images. We only used the image where all the three models correctly predict and agree on the label. From a million images that we originally collected using the Flickr API, only 434,337 (43%) remained after this filtering process. The fact that different models are able to correctly classify the same image means that the image is robust enough when there is no perturbation (Easy Images). From the three different attacks used, we decided to use PGD attack for the rest of the experiments given that it is the strongest attack as shown in Table 2. Moreover, we use ResUpNet and BaRT defenses to find the defense-friendly images.

For the easy images subset, a PGD attack was implemented per image using  $\epsilon$  values of 0.01, 0.02 and 0.04 for each model. We defined those adversarial images as  $\epsilon$ -robust if they were still classified correctly by the three models without using any defense algorithm. Depending on the amount of perturbation, and the ability of the image to retain their accuracy despite the added perturbation, some images are more robust than others. From the 434,337 robust images collected, only 15,554 images were able to retain the original prediction across different classifiers for adversarial attacks with an  $\epsilon = 0.01$ . Those images were used to create the new robust dataset.

With extensive experimentation, we found that those images that are more *robust* (can stand greater levels of perturbation in an attack) are also in general more *defense-friendly*. Therefore, we consider defense-friendly images as a special case of robust images which makes them easier to recover from an attack. To better understand the difficulty of a dataset in regards to the attacks and defenses used, we introduced three dataset metrics that are further explained in Section 5.3. Figure 1 shows the frequency of classes (according to the ImageNet labels), which has the greater number of correct samples in descending order with at least 30 occurrences. For this robust subset, the mean frequency of a the classes is 25.75 and the median 7. There are

604 classes with at least one image and the remaining 396 classes contain no image at all. A random sample of images corresponding to the three categories from our ImageNet-R dataset can be seen in Figure 2.

## 5. Performance Evaluating

Table 2: Accuracy of datasets with no defense: Datasets (D 1-3) get very similar results with very low standard deviation between them. However, datasets ‘Robust’ (R 1-3) consistently get better results than the randomly selected datasets. Values given in percentage and all datasets have a total of 5,000 images.

Model	$\epsilon=0$	FGSM			PGD			DDN			
		$\epsilon=0.01$	$\epsilon=0.02$	$\epsilon=0.04$	$\epsilon=0.01$	$\epsilon=0.02$	$\epsilon=0.04$	$n=20$	$n=40$	$n=60$	
D1	vgg	71.8	4.1	4.2	4.6	0.0	0.0	0.2	1.4	0.9	0.8
	resnet	75.9	7.4	7.6	9.6	0.0	0.0	0.0	0.4	0.2	0.1
	dense	74.6	4.2	4.3	5.6	0.4	0.2	0.0	0.3	0.1	0.1
D2	vgg	70.0	4.0	3.7	4.5	0.5	0.3	0.2	1.3	0.1	0.1
	resnet	74.8	6.7	6.3	8.0	0.0	0.0	0.0	0.3	0.1	0.1
	dense	73.0	3.4	3.4	4.9	0.1	0.0	0.0	0.2	0.8	0.1
D3	vgg	71.8	4.7	4.3	4.8	0.7	0.5	0.2	1.4	1.1	1.0
	resnet	76.3	7.8	7.6	9.4	0.1	0.0	0.0	0.6	0.3	0.2
	dense	74.6	4.0	3.7	5.4	0.1	0.0	0.0	0.3	0.2	0.1
R1	vgg	99.2	83.6	82.5	79.2	54.1	26.1	10.9	83.2	83.2	83.2
	resnet	100	100	100	100	99.8	45.0	17.7	100	100	100
	dense	100	100	100	100	100	54.6	22.9	100	100	100
R2	vgg	99.3	83.6	82.3	79.8	54.1	26.2	11.2	83.6	83.6	83.6
	resnet	100	100	100	100	99.8	44.0	17.8	100	100	100
	dense	100	100	100	100	100	54.9	23.2	100	100	100
R3	vgg	99.3	83.6	82.3	79.8	54.1	26.2	11.2	83.6	83.6	83.6
	resnet	100	100	100	100	99.8	44.0	17.8	100	100	100
	dense	100	100	100	100	100	54.9	23.2	100	100	100

### 5.1. Evaluation for Adversarial Attacks

The idea in this section is to compare how the accuracy of different models is affected when using different attacks on a “normal” dataset (randomized subset from ImageNet) and the new dataset ImageNet-R. In this section, we present the results from our experiments which further support our claim that some images are more difficult to perturb or more robust. We observe a significant difference in the accuracy between the new dataset and a “normal” dataset under different attacks for ImageNet-based models. We randomly split robust images into three groups R (1-3). The same process was done for the normal datasets D (1-3). Table 2 shows that robust datasets have more images with 100% or nearly 100% accuracy on non-perturbed images, and they are also much more resilient to attacks for small perturbations  $\epsilon$  in comparison to randomly selected datasets.

### 5.2. Evaluation for Adversarial Defenses

Similar to the evaluation of adversarial attacks, a random dataset and the easy images (including Robust Images and Defense-friendly) dataset will be compared. If our hypothesis and claim hold true, those images that are more robust or difficult to misclassify should also be more defense-friendly or easier to recover when using a defense.

As can be seen in Table 1, defense-friendly images are more resilient to attacks and they also recover better when using a defense. It also shows that for non-defense friendly

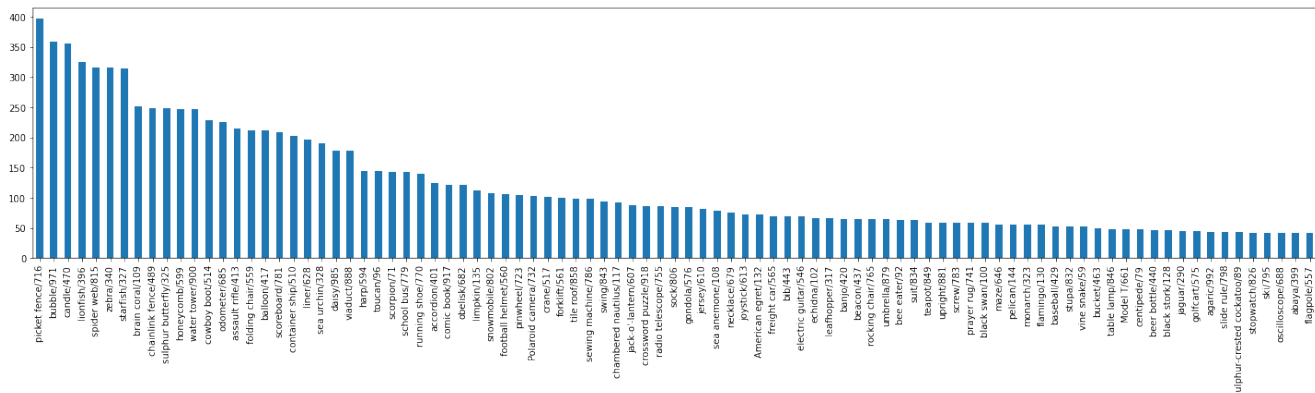


Figure 1: This figure shows the per class frequency from the newly collected dataset (15,554 images) with the label and index number according to the ImageNet dataset. The figure shows only the classes that have at least 30 occurrences arranged in descending order.



Figure 2: Sample of images from our curated dataset ImageNet-R for the categories: Easy,  $\epsilon$ -robust and Defense-friendly images.

images, the percentage of accuracy recovery while using a defense is much lower than a more robust dataset.

## 6. Metrics for Quantifying Dataset Robustness

To measure the overall robustness of a dataset, we propose three different metrics, namely, ARD, AMP and ADF. A general convention has been to use the difference between the top 2 predictions or logit scores as a confidence metric which is not reliable in general. We validate this by comparing the  $\ell_2$  difference between the top 2 logit scores on ResNet50 for a pair of robust and non-robust datasets. Figure 3 shows our results. We can see that many non-robust images have a greater distance than that of robust images. On the other hand, our method uses a combination of three metrics that are more reliable and are designed to measure possible bias globally in a dataset as a whole rather than locally at individual images.

**Adversarial Robust Dataset (ARD) Score:** This score provides the proportion of images in a dataset (range [0, 1]) that are robust given an attack  $atk$  with a perturbation  $\epsilon$  and a model  $M$ .

**Adversarial Minimum Perturbation (AMP) score:** This metric computes the minimum  $\epsilon$  perturbation in which a proportion of the images  $P$  in a dataset  $D$  are not able to resist an specific attack  $atk$  for the model  $M$ . The higher the  $\epsilon$  value returned by the function  $AMP$ , the easier the image dataset is to defend.

**Adversarial Defense-friendly (ADF) Score:** This score finds the proportion of images in a dataset that are recov-

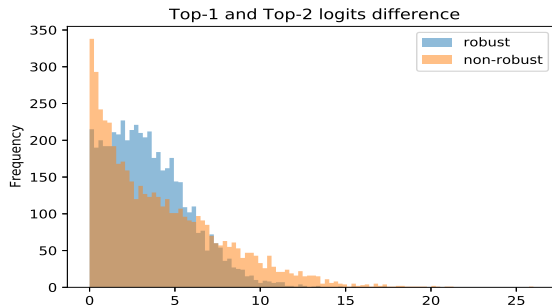


Figure 3: Distribution of the  $\ell_2$  distance between the logit scores for the top-1 and top-2 predictions for a pre-trained ImageNet model (ResNet50) evaluated on a robust dataset and a randomized non-robust dataset. The third color in the figure shows their intersection. The difference between top-1 and top-2 logits is not always greater in robust images. In many cases it can also happen that the  $\ell_2$  distance in the logits is similar in robust and non-robust.

ered using a defense for a specific attack with a small  $\epsilon$ . We created 6 different subsets of images (robust and non-robust) with images randomly selected from the collected dataset (Section 4). The results from the ARD, AMP and ADR Score applied to those datasets are shown in Table 3. Despite some models performing better, in general, robust images have a higher ADF score. We can also observe in the results that there are some transferability of robustness when we test on different models.

Initially, to differentiate robust images from non-robust ones, we use three losses (corresponding to the three models

### Algorithm 1: ARD Score

---

**Input:** An adversarial attack  $atk$ , perturbation  $\epsilon$ , dataset  $D$ , model  $M$   
**Output:** ARD Score denoted by  $score$   
**CalculateARDScore** ( $atk, \epsilon, D, M$ )

```

R ← [] // Empty list for Robust Images
// For each image i in dataset D
foreach i ∈ D do
    // Generate Adversarial Image
    iadv ← atk(i, ε, M)
    if M(i) = M(iadv) then
        R.insert(i)
score = Length(R) / Length(D)
return score

```

---

### Algorithm 2: AMP Score

---

**Input:** An adversarial attack  $atk$ , perturbation  $\alpha$ , dataset  $D$ , model  $M$ , threshold  $P$   
**Output:** AMP Score denoted by  $\epsilon$   
**CalculateAMPScore** ( $atk, \epsilon, D, M, P$ )

```

!R ← []
foundε ← False
ε ← 0 // Initialize ε as zero
while foundε = False do
    // For each image i in dataset D
    foreach i ∈ D do
        // Generate Adversarial Image
        iadv ← atk(i, ε, M)
        if M(i) ≠ M(iadv) then
            !R.insert(i)
A ← Length(!R) / Length(D)
if A ≥ P then
    foundε ← True
else
    ε ← ε + α
return ε

```

---

### Algorithm 3: ADF Score

---

**Input:** An adversarial attack  $atk$ , perturbation  $\epsilon$ , dataset  $D$ , model  $M$   
**Output:** ADF Score denoted by  $score$   
**CalculateADFScore** ( $atk, \epsilon, D, M$ )

```

// Empty list of Defense-friendly Images
DF ← []
// For each image i in dataset D
foreach i ∈ D do
    // Generate Adversarial Image
    iadv ← atk(i, ε, M)
    if M(i) = M(iadv) then
        DF.insert(i)
score = Length(DF) / Length(D)
return score

```

---

DenseNet, ResNet and Inception) and an epsilon value as a parameter to calculate the adversarial perturbations for that image. If the adversarial image resists a small  $\epsilon$  value, we consider the image to be  $\epsilon$ -robust. However, this method is computationally expensive and hence, there is a need for a metric that can be computed more efficiently. One plausible

Table 3: ARD, AMP and ADF scores for 6 different datasets, each containing 5,000 randomly selected images from our curated dataset. We used PGD attack for this experiment. Robust datasets (R1-R3) obtain better scores for each metric while Non-robust (NR1-NR3) datasets have significantly lower scores.

Data	Models	ARD Score			AMP Score				ADF Score					
		$\epsilon$			Threshold				BaRT		ResUpNet			
		$\epsilon$	$\epsilon$	$\epsilon$	50	70	80	90	$\epsilon$	$\epsilon$	$\epsilon$	$\epsilon$		
NR1	vgg	0.03	0.01	0.01	0.00	0.01	0.01	0.01	0.06	0.04	0.02	0.16	0.14	0.10
	resnet	0.04	0.01	0.00	0.00	0.01	0.01	0.01	0.11	0.07	0.04	0.35	0.31	0.27
	dense	0.07	0.01	0.00	0.01	0.01	0.01	0.01	0.15	0.08	0.04	0.38	0.31	0.25
NR2	vgg	0.03	0.01	0.01	0.00	0.01	0.01	0.01	0.08	0.05	0.02	0.27	0.23	0.19
	resnet	0.03	0.00	0.00	0.01	0.01	0.01	0.01	0.12	0.07	0.04	0.33	0.29	0.24
	dense	0.08	0.02	0.00	0.01	0.01	0.01	0.01	0.18	0.09	0.07	0.41	0.34	0.27
NR3	vgg	0.04	0.01	0.01	0.00	0.01	0.01	0.01	0.07	0.04	0.03	0.27	0.24	0.20
	resnet	0.03	0.00	0.00	0.00	0.01	0.01	0.01	0.11	0.06	0.03	0.22	0.18	0.14
	dense	0.07	0.01	0.00	0.01	0.01	0.01	0.01	0.17	0.10	0.06	0.42	0.35	0.27
R1	vgg	0.51	0.25	0.11	0.01	0.01	0.01	0.01	0.56	0.36	0.22	0.66	0.63	0.57
	resnet	1.00	0.45	0.18	0.02	0.02	0.02	0.03	0.98	0.62	0.37	0.85	0.82	0.77
	dense	1.00	0.55	0.23	0.02	0.03	0.03	0.03	0.98	0.70	0.45	0.89	0.87	0.84
R2	vgg	0.53	0.25	0.11	0.01	0.01	0.01	0.01	0.59	0.36	0.21	0.68	0.66	0.59
	resnet	1.00	0.45	0.19	0.02	0.02	0.02	0.03	0.98	0.63	0.39	0.85	0.84	0.78
	dense	1.00	0.55	0.22	0.02	0.03	0.03	0.03	0.99	0.71	0.42	0.91	0.89	0.83
R3	vgg	0.53	0.26	0.11	0.01	0.01	0.01	0.01	0.57	0.36	0.21	0.68	0.65	0.59
	resnet	1.00	0.44	0.18	0.02	0.02	0.02	0.02	0.97	0.59	0.37	0.87	0.83	0.79
	dense	1.00	0.57	0.23	0.02	0.03	0.03	0.03	0.99	0.70	0.43	0.91	0.90	0.84

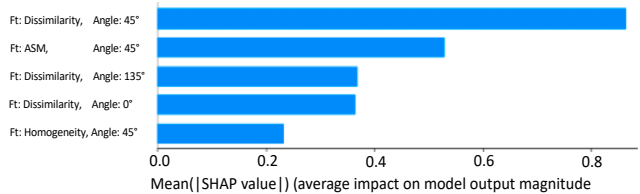


Figure 4: This figure shows a summary SHAP plot displaying the top 5 most important features (extracted using GLCM for the best performing ML model in Table 4) and their contribution for every sample. This plot sorts features by the sum of SHAP value magnitudes over all samples. In addition, it shows the distribution of the impacts of every features with respect to the output. As a first insight, the model finds more useful features from distance 1 given in pixels. Moreover, the most significant texture feature seems to be dissimilarity (3 out of 5 features in the top 5 features) followed by the ASM feature using angle of 45°. A higher value of dissimilarity indicates that an image is more likely to be a non-robust image. On the contrary, the lower the dissimilarity in the image the higher the likelihood of being a robust image.

solution is to use the first metric as the ‘ground truth’ labels and train a second classifier to classify the image as ‘robust’ or ‘non-robust’. The robust class is relative to a minimum level of perturbation  $\epsilon$  which must be provided to create the ‘ground truth’ labels.

We used the GLCM features (see Section 3.3) to extract statistical features from the 15,554 images that are considered robust and then randomly selected the same number of images from the non-robust pool to get a balanced dataset. A binary classifier (robust vs non-robust) was trained with the statistical features extracted using GLCM [34]. Figure 4 illustrates the importance of the statistical features such as dissimilarity in that binary classifier. ImageNet-trained CNNs are biased towards texture [11], therefore, we want to analyse texture and determine if it is possible to extract

Table 4: Robust vs Non-robust classification of images. Before classifying the actual contents of an image, these classifiers determine whether an image is robust or not i.e. whether it is likely to be correctly classified by a deep model or not. GLCM features were extracted from the Y-channel of the YCbCr color space [28].

Model	Type Feature	Accuracy	Precision	Recall
Logistic Regression	Statistical GLCM	66.1	69.6	54.9
Random Forest	Statistical GLCM	75.0	74.6	74.5
SVM	Statistical GLCM	74.0	73.9	72.9
ResNet34 (Finetuned)	Deep Learning / CNN	87.9	83.0	92.1
ResNet152 (Finetuned)	Deep Learning / CNN	88.6	85.3	91.3
DenseNet121 (Finetuned)	Deep Learning / CNN	88.0	84.1	92.2

texture elements that could indicate when an image is more robust than others. Next, a CNN model is trained in the classification problem (robust vs non-robust) to compare the results of the statistical approach versus the Deep Learning approach.

For the statistical features, we train a simple Logistic Regression (baseline), a Random Forest algorithm and a Support Vector Machine (SVM). For every image in the training dataset we extract the GLCM properties including *contrast*, *dissimilarity*, *homogeneity*, *ASM*, *energy* and correlation for 4 different angles (rotationally invariant features)  $0, 45, 90, 135$ . Additionally, to calculate those properties it is necessary to provide the distance given in pixels. In this experiment, we use distances of 1, 2 and 3 pixels. After calculating those features, we have a total of 72 features (6 properties x 4 angles x 3 distances = 72 features). In the statistical approach, it is clear that elements such as color or other important features of the images will be missing. However, if texture is important for CNN classifiers in general, those same texture features might be sufficient to determine how well an image is recognizable by a classifier and if they are more resilient to attacks than other randomly selected images.

Table 4 shows the results of traditional machine learning (ML) approaches and deep learning models trained to determine whether an image is robust or non-robust. For the non-deep-learning approaches, 72 features are extracted using the statistical method GLCM. Despite the expected results of deep learning models performing better than traditional ML, the best performing ML model achieves 75 accuracy using only statistical features extracted from a gray-scale image. For the gray-scale image we used the Y channel from the YCbCr. According to Pestana et al. [28], the Y channel from YCbCr color space has more relevant features than the color CbCr channels. In addition, Geirhos et al. [11] discuss the importance of texture for CNN ImageNet-based classifiers. The results from Table 4 show that a model using extracted GLCM features from the Y-channel is able to recognise in the majority of cases whether an image is robust or not. An interesting point to note is that while deep models achieve higher accuracy using RGB images, the results in Table 4 also confirm that the most relevant information comes from the Y-channel and

Table 5: Robust vs Non-robust classification of images under attack. The same classifiers from Table 4 are tested using adversarial images from a PGD attack  $\epsilon=0.01$ . Despite the initial accuracy of deep learning models, they seem to be much more susceptible than traditional ML classifiers using statistical GLCM features.

Model	Type Feature	Accuracy	Precision	Recall
Logistic Regression	Statistical GLCM	63.70	72.2	52.80
Random Forest	Statistical GLCM	73.08	75.80	77.30
SVM	Statistical GLCM	68.08	79.10	70.10
ResNet34 (Finetuned)	Deep Learning / CNN	78.90	70.43	99.60
ResNet152 (Finetuned)	Deep Learning / CNN	67.07	97.40	75.30
DenseNet121 (Finetuned)	Deep Learning / CNN	67.29	97.00	75.70

texture of an image. In Table 5, we add PGD perturbations with  $\epsilon = 0.01$  to the images and repeat the experiment of Table 4. Our results show that the accuracy of deep models drops more drastically compared to the shallow models that use only statistical GLCM features from gray images for the same attack and amount of perturbation. This confirms that our simple feature based classifiers are good proxies to measure the robustness of a dataset.

## 7. Conclusion

We showed the existence of *defense-friendly* images that are both resilient (or robust) to adversarial attacks and also recover the under attack model’s accuracy more easily than other images. We demonstrated that the classification accuracy of models on images that are initially classified correctly by models of different architectures (but trained using the same ImageNet dataset) is easier to recover using a defense. This phenomenon may give the misleading impression regarding the performance of a given defense mechanism and has been exploited by some adversarial defense works which hand-pick a subset of clean images to report a near 100 defense accuracy. However, Akhtar et al. [1] argue that evaluating defense mechanisms on already misclassified images is not meaningful and such images should not be considered in evaluation since an attack on a misclassified image is considered successful by default and this could mislead the interpretation of the results. We showed through extensive experimentation that such hand-picked datasets indeed give misleading advantages while evaluating the performance of a defense. To overcome this problem, we proposed three metrics to quantify the robustness/defense-friendliness of a dataset. We also provide a dataset with more than 15k robust images to complement the 7.5K natural adversarial examples of [17]. We believe that our dataset and metrics will be valuable for unbiased benchmarking of defense mechanisms.

## 8. Acknowledgements

The main author was recipient of an Australian Government Research Training Program (RTP) Scholarship at The University of Western Australia. This research was supported by ARC Discovery Grant DP190102443. We would like to thank Robyn Owens for her insightful comments.



## References

- [1] Naveed Akhtar, Jian Liu, and Ajmal Mian. Defense against universal adversarial perturbations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3389–3398, 2018.
- [2] Naveed Akhtar and Ajmal Mian. Threat of adversarial attacks on deep learning in computer vision: A survey. *arXiv preprint arXiv:1801.00553*, 2018.
- [3] Arjun Nitin Bhagoji, Daniel Cullina, and Prateek Mittal. Dimensionality reduction as a defense against evasion attacks on machine learning classifiers. *arXiv preprint arXiv:1704.02654*, 2017.
- [4] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.
- [5] Jeremy M Cohen, Elan Rosenfeld, and J Zico Kolter. Certified adversarial robustness via randomized smoothing. *arXiv preprint arXiv:1902.02918*, 2019.
- [6] Nilaksh Das, Madhuri Shanbhogue, Shang-Tse Chen, Fred Hohman, Li Chen, Michael E Kounavis, and Duen Horng Chau. Keeping the bad guys out: Protecting and vaccinating deep learning with jpeg compression. *arXiv preprint arXiv:1705.02900*, 2017.
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [8] Gavin Weiguang Ding, Luyu Wang, and Xiaomeng Jin. AdverTorch v0.1: An adversarial robustness toolbox based on pytorch. *arXiv preprint arXiv:1902.07623*, 2019.
- [9] Gintare Karolina Dziugaite, Zoubin Ghahramani, and Daniel M Roy. A study of the effect of jpg compression on adversarial images. *arXiv preprint arXiv:1608.00853*, 2016.
- [10] N Benjamin Erichson, Zhewei Yao, and Michael W Mahoney. Jumprelu: A retrofit defense strategy for adversarial attacks. *arXiv preprint arXiv:1904.03750*, 2019.
- [11] 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. *arXiv:1811.12231 [cs, q-bio, stat]*, Jan. 2019. arXiv: 1811.12231.
- [12] Ian J Goodfellow et al. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [13] Chuan Guo, Mayank Rana, Moustapha Cisse, and Laurens Van Der Maaten. Countering adversarial images using input transformations. *arXiv preprint arXiv:1711.00117*, 2017.
- [14] Robert M Haralick, Karthikeyan Shanmugam, and Its' Hak Dinstein. Textural features for image classification. *IEEE Transactions on systems, man, and cybernetics*, (6):610–621, 1973.
- [15] 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*, pages 770–778, 2016.
- [16] Warren He, James Wei, Xinyun Chen, Nicholas Carlini, and Dawn Song. Adversarial example defense: Ensembles of weak defenses are not strong. In *11th {USENIX} Workshop on Offensive Technologies ({WOOT} 17)*, 2017.
- [17] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples, 2019.
- [18] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [19] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, Cite-seer, 2009.
- [20] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [21] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [22] Xuanqing Liu and Cho-Jui Hsieh. Rob-gan: Generator, discriminator, and adversarial attacker. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11234–11243, 2019.
- [23] Chunchuan Lyu, Kaizhu Huang, and Hai-Ning Liang. A unified gradient regularization family for adversarial examples. In *2015 IEEE International Conference on Data Mining*, pages 301–309. IEEE, 2015.
- [24] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [25] Nina Narodytska and Shiva Prasad Kasiviswanathan. Simple black-box adversarial perturbations for deep networks. *arXiv preprint arXiv:1612.06299*, 2016.
- [26] Nicolas Papernot et al. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, pages 506–519. ACM, 2017.
- [27] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [28] Camilo Pestana, Naveed Akhtar, Wei Liu, David Glance, and Ajmal Mian. Adversarial perturbations prevail in the y-channel of the ycbcr color space, 2020.
- [29] Edward Raff, Jared Sylvester, Steven Forsyth, and Mark McLean. Barrage of random transforms for adversarially robust defense. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6528–6537, 2019.
- [30] Jerome Rony, Luiz G. Hafemann, Luiz S. Oliveira, Ismail Ben Ayed, Robert Sabourin, and Eric Granger. Decoupling direction and norm for efficient gradient-based l2 adversarial attacks and defenses. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [31] Jérôme Rony, Luiz G Hafemann, Luiz S Oliveira, Ismail Ben Ayed, Robert Sabourin, and Eric Granger. Decoupling direc-

- tion and norm for efficient gradient-based  $l_2$  adversarial attacks and defenses. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4322–4330, 2019.
- [32] Uri Shaham, Yutaro Yamada, and Sahand Negahban. Understanding adversarial training: Increasing local stability of supervised models through robust optimization. *Neurocomputing*, 307:195–204, 2018.
- [33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [34] S. Singh, D. Srivastava, and S. Agarwal. Glcm and its application in pattern recognition. In *2017 5th International Symposium on Computational and Business Intelligence (IS-CBI)*, pages 20–25, 2017.
- [35] David Stutz, Matthias Hein, and Bernt Schiele. Disentangling adversarial robustness and generalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6976–6987, 2019.
- [36] Christian Szegedy et al. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [37] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [38] Saeid Asgari Taghanaki, Kumar Abhishek, Shekoofeh Azizi, and Ghassan Hamarneh. A kernelized manifold mapping to diminish the effect of adversarial perturbations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11340–11349, 2019.
- [39] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*, 2017.
- [40] Cihang Xie, Yuxin Wu, Laurens van der Maaten, Alan L Yuille, and Kaiming He. Feature denoising for improving adversarial robustness. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 501–509, 2019.
- [41] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. *arXiv preprint arXiv:1704.01155*, 2017.